Top-Down Structure Influences Learning of Non-Adjacent Dependencies in an Artificial Language

Felix Hao Wang, Jason D. Zevin, and Toby H. Mintz

University of Southern California

Author Note

Felix Hao Wang, Department of Psychology, University of Southern California. Jason D. Zevin, Department of Psychology and Department of Linguistics, University of Southern California. Toben H. Mintz, Department of Psychology and Department of Linguistics, University of Southern California.

Felix Wang is now at Department of Psychology, University of Pennsylvania.

Correspondence concerning this article should be addressed to Felix Hao Wang, 425 S University Ave, Philadelphia, PA 19104. Contact: haowang1@sas.upenn.edu.

Portions of the work were presented at the 38th annual meeting of the Cognitive Science Society in Philadelphia, PA.

Abstract

Due to the hierarchical organization of natural languages, words that are syntactically related are not always linearly adjacent. For example, the subject and verb in *the child always runs* agree in person and number, although they are not adjacent in the sequences of words. Since such dependencies are indicative of abstract linguistic structure, it is of significant theoretical interest how these relationships are acquired by language learners. Most experiments that investigate non-adjacent dependency (NAD) learning have used artificial languages in which the to-be-learned dependencies are isolated, by presenting the minimal sequences that contain the dependent elements. However, dependencies in natural language are not typically isolated in this way. We report the first demonstration to our knowledge of successful learning of *embedded* NADs, in which silences do not mark dependency boundaries. Subjects heard passages of English with a predictable structure, interspersed with passages of the artificial language. The English sentences were designed to induce boundaries in the artificial languages. In Experiment 1 & 3 the artificial NADs were contained within the induced boundaries and subjects learned them, whereas in Experiment 2 & 4, the NADs crossed the induced boundaries and subjects did not learn them. We take this as evidence that sentential structure was “carried over” from the English sentences and used to organize the artificial language. This approach provides several new insights into the basic mechanisms of NAD learning in particular and statistical learning in general.

Keywords: Language acquisition; Bracketing; Grammatical Entrainment; Statistical learning; Non-adjacent dependency

Due to the hierarchical organization of the syntax of natural languages, lexical items (and morphemes) that are syntactically related are not always linearly adjacent. Thus, to acquire the specifics of the hierarchical grammar, learners must be able to track both adjacent and non-adjacent dependencies in a linear sequence of words. For example, in *the child runs*, the third person singular subject, *child*, and the agreeing inflected verb, *runs*, are linearly adjacent; however, they are non-adjacent in *the child always runs*. For language learners just beginning to learn their language's syntax, evidence about which elements are related in grammatical processes could provide extremely useful information for further grammatical learning. This kind of information is not only beneficial for theories that view language acquisition as a largely domain-general learning problem, but also for theories in which learners have innate domain-specific constraints on representing and processing language. Therefore, an important question in language acquisition research is how learners detect adjacent and non-adjacent grammatical dependencies, as doing so could help learners understand how their language is structured.

There has been considerable interest in investigating learning mechanisms that could detect dependencies in linear sequences within spoken utterances. Many accounts of adjacent dependency learning posit that learners track statistical information about the transitions from one unit (e.g., syllable or word) to the next (Romberg & Saffran, 2013; Saffran, Aslin, & Newport, 1996). However, while most studies of adjacent dependencies have found success in a variety of learning scenarios (Pelucchi, Hay, & Saffran, 2009; Saffran et al., 1996; Thompson & Newport, 2007), the studies of non-adjacent dependencies to date have only found learning in limited situations, with some studies reporting success in learning and others reporting failure. (e.g., success: Gómez, 2002; Gómez & Maye, 2005; Romberg & Saffran, 2013; failure: Newport & Aslin, 2004; mixed: Peña, Bonatti, Nespor & Mehler, 2002; Santelmann & Jusczyk, 1998). In

those experiments that showed successful learning of non-adjacent dependencies with artificial languages, the minimal sequences that contained the dependencies were presented as discrete, short, pre-segmented sub-sequences. Specifically, the sub-sequences were surrounded by silences, and the (non-adjacent) dependent elements were found at sub-sequence edges. In experiments investigating non-adjacent dependencies between words, trigrams were composed of words in which the third word was dependent on the first word. Participants were presented the trigrams one at a time, with silence intervening between presentations, as in (1a), where pause length is indicated by the number of pound signs and elements making up dependencies are bolded (Gómez, 2002; Gómez, Bootzin & Nadel, 2006; Gómez & Maye, 2005; Romberg & Saffran, 2013; Wang & Mintz, under review). Similarly, in experiments investigating non-adjacent dependencies between syllables in syllable sequences, learning occurred only when brief silences were introduced before each trigram composed of syllables, as in (1b) (Peña et al., 2002). When syllables were concatenated continuously, participants showed no learning (see also Newport & Aslin, 2004)¹. Moreover, Wang and Mintz (under review) embedded sub-sequences into longer ‘sentences’ sequences (in 1c) with a similar structure to those in Gómez (2002), and found that such embedding hindered adults’ ability to learn the dependencies.

¹ Newport and Aslin (2004) also reported apparent non-adjacent dependency learning between non-adjacent segments in adjacent syllables, where the segments matched in type (i.e., vowel and vowel, consonant and consonant). Newport and Aslin explained their findings in the framework of a prominent phonological theory, in which consonants and vowels reside on separate representational tiers; on those tiers, the dependencies subjects learned where, in fact, adjacent.

- (1) a. ## pel # wadim # rud ## vot # kicey # jic ## ...
 b. # pu ra ki # be li ga # ta fo gu # pu fo ki # ...
 c. ## blit # choon # balip # jub # ghire # paziv ##

Thus, in the studies just discussed, subjects' success in learning non-adjacent dependencies apparently depended on whether the sequences were pre-segmented. Why might this be? One possibility is that pre-segmenting the material in this way brackets the sequence such that both dependent elements are in an edge position. Peña et al. (2002) propose that the mechanisms involved in learning adjacent and non-adjacent dependencies are different, in part because of the role of edges in computing the latter, but not the former (see also Endress & Bonatti, 2007). Specifically, they propose that adjacent dependencies are learned as a result of statistical processing of co-occurrence patterns, whereas non-adjacent patterns are learned by forming rules that capture patterns relating material at the edges of sequences. Indeed, Endress, Nespor & Mehler (2009) argued that edges are privileged in the kind of position-related computations they afford, and that placement at edges could be an important constraint for learning non-adjacent dependencies.

However, a different and perhaps simpler explanation is that the same learning mechanisms are involved in both kinds of pattern learning, but non-adjacent dependency learning is less robust because the distance between dependent elements puts greater demands on the processing and remembering of the relevant parts of the sequence. In this case, silences may simply serve as a way of breaking a sequence down into a more manageable size (Wang & Mintz, under review). In addition, when dependent elements occur at the beginnings and ends of sequences (i.e., edges), learning might be facilitated simply because those positions are more cognitively salient (Slobin, 1973).

We will return to the issue of how silences may benefit non-adjacent dependency learning in a moment. Whatever the explanation, in natural languages non-adjacent dependencies are not restricted to positions before or after silences, and are often embedded in longer sequences. This may pose a problem for learners acquiring such dependencies, if learning non-adjacent dependencies critically requires the presence of silences. In that case, only dependencies that occur at utterance boundaries would be learnable. On the other hand, whereas silences provide boundaries and edges in the input signal, it is possible that boundaries that arise as a result of listeners processing the input stream could also facilitate dependency learning. That is, perhaps boundaries that listeners *compute* as the result of syntactic parsing play a similar bracketing role to boundaries that occur in the speech signal as the result of silences. Under this interpretation, silences per se might not be critical for learning non-adjacent dependencies, but rather any kind of boundary could facilitate the detection of non-adjacent dependencies when the elements making up the dependency are at edge initial or final positions. This possibility is consistent with the general idea that edges are important for detecting dependencies (Endress & Bonatti, 2007; Endress et al., 2009; Peña et al., 2002), but takes a broader view on what constitutes an edge. However, it is also consistent with a more general account whereby the role of edges (however derived) may be to guide learners to the locations in the longer sequence that contain the dependencies. On this account, the necessity of edges in learning non-adjacent dependencies is not due to the constraints of special processing mechanisms, but rather it arises because non-adjacency renders dependencies more difficult to detect, and thus they benefit from more guided, focused processing.

In this paper we investigate this question. We ask whether edges or boundaries that are *computed* from syntactic parsing function similarly to silences in the signal. This would be a

theoretically important finding, as it would demonstrate that detecting statistical patterns in quasi-linguistic input is guided by abstract, top-down knowledge. To test this empirically, we needed a method of inducing bracketing via syntactic parsing that does not involve silences. To this end, we adopt a recently described technique from Ding, Melloni, Zhang, Tian and Poeppel (2016), which we call grammatical entrainment. Sentences with identical syntactic structures are presented in a continuous sequence of words, cyclically repeating the grammatical structure: e.g. *...dry fur rubs skin fat rat sensed fear new plans gave hope....* Ding et al. found that this kind of syntactic repetition results in neurophysiological signatures that demonstrate that listeners track the syntactic structures, including the sentence boundaries every four words, despite the absence of silences between utterances (see also Luo & Poeppel, 2007; Wang, Ding, Ahmar, Xiang, Poeppel, & Simon, 2012). We reasoned that learners might impose entrained structural processing onto an artificial language that is presented immediately following the cyclic presentation of the English structures. In other words, we hypothesized that subjects' repeated structure building of familiar material (English) would guide their parsing of novel material (the artificial language). Learners would thereby impose a syntactic bracketing onto the artificial language, resulting in a sequence of four-word "sentences", with abstract boundaries between them. If learners then acquire the dependencies in the artificial language, it demonstrates that silences, *per se*, are not critical for NAD learning, but rather top-down bracketing information can play a similar role. Given the same periodicity between English sentences and the artificial language sentences, words in the artificial language would have *phase* assigned to them and the dependencies in the artificial stream could be anchored to the repeating English structures. Specifically, the entrainment method would then allow us to manipulate whether or not the sub-sequences contain the NADs, or whether the bracketing disrupts the NADs, without changing the

actual artificial language sequence in any way. Because this manipulation changes the alignment of the cyclic English structure and the dependencies in the artificial language, we call this a phase manipulation.

Experiment 1

The goal of Experiment 1 was to determine whether subjects would learn non-adjacent dependencies in continuous sequences of words in an artificial language, when the sequences were preceded by sequences of English words that could be parsed into a series of four-word sentences. We hypothesized that the the sentence level structures that listeners generated when parsing the English words would force an analogous parsing in the artificial language. We designed the sequences so that the four-word sequences in the artificial language each contained a non-adjacent dependency. That is, the grammatical entrainment of English was *in phase* with the dependencies in the artificial language.

Methods

Participants. Twenty-four undergraduate students at University of Southern California were recruited from Psychology Department subject pool. Half of them participated in each counterbalancing condition. The sample size was based on previous studies (Newport & Aslin 2004).

Stimuli and Design. We recorded speech from a native English speaker and digitized the recording at a rate of 44.1 kHz. We recorded 2 types of words: English words and nonsense words. For English words, we recorded 5 names (*Brian, John, Kate, Nate, Clair*), 5 monosyllabic verbs in 3rd person singular form (*turns, keeps, puts, lets, has*), 5 pronouns (*these, those, this, that, it*) and 5 particles (*down, on, up, off, in*). For the non-adjacent dependency, we used 12 novel words for 4 different positions (YAXB): 3 *Y* words at position 1 (*pel, tink, blit*), 3 *A* words

at position 2 (*swech*, *voy*, *rud*), 3 *X* words at position 3 (*dap*, *tood*, *wesh*) and 3 *B* words at position 4 (*ghire*, *jub*, *tiv*).

We then spliced the words into individual sound files that ranged from 300ms to 737ms in duration. We used temporal scaling features in Praat (Boersma, 2001) to compress all the words to 250ms. An additional 83ms of silence was added to the end of *each* word to increase intelligibility, and to ensure that when words were concatenated in a continuous stream for presentation, they occurred at 3Hz (Ding et al., 2015).

Training materials. To create the training stream, we interleaved English sentences and artificial sentences together. A total of 432 ‘sentences’ made of non-words and 858 English sentences were concatenated together in the following manner. English sentences all had a [Name Verb Pronoun Particle] structure, where each English sentence was created by randomly picking a name, a verb, a pronoun, and an particle, in that order. As such, each sentence consisted of 4 syllables (with the exception of sentences containing the word *Brian*), and lasted 1.33 seconds. Since words were randomly selected and constrained only by position, there were no statistically reliable dependencies between particular words in the English portion of the training materials².

Each sentence in the artificial language was a concatenation of 4 non-words. As with the English sentences, each sequential position in a sentence was filled by one of three possible words, and the set of possible words was different for each position. We represent this pattern

² While we control for statistical dependencies in the stimuli, there are likely statistical dependencies in the subjects’ prior knowledge of English that could influence processing of the experimental material. We take up this issue again in Experiment 3.

as YAXB, with one letter for each class of word in the artificial language. However, whereas in the English sentences there were no statistically reliable dependencies across positions, in the artificial sentences, the second word perfectly predicted the fourth word, and vice-versa; we depict this as YA_iXB_i . Given there were 3 different words for each non-dependent position (i.e., positions 1, 2, and 3), there were 27 possible artificial sentences, each of which occurred 16 times, resulting in the 432 artificial language sentences.

The training stream was constructed by concatenating alternating sequences of English sentences and sentences in the artificial language. First, we concatenated 5-7 English sentences to create the entrainment effect. After the English sentences, three artificial language sentences were presented, followed by another set of 5-7 English sentences (Figure 1). Aside from the 83 ms of silence between *all* words, there were no additional silences between words or sentences of English or the artificial language.

Figure 1. The design of language materials in the training phase of Experiment 1. The English sentence and the artificial dependencies are in phase.

Figure 1 depicts the grammatical bracketing of three English sentences and the entrained bracketing of one artificial sentence in sequence. Note that, in addition to the sentence boundaries and dependency in the artificial language, Figure 1 also indicates a grammatical

dependency in the English sentences between the verb and the particle (e.g., *puts_down*).

Although this similarity means that the English and artificial sentences are parallel at a more fine-grained level than the sentence-level bracketing, our claims in this paper focus on alignment at the level of the sentence, although we also address this issue further in Experiment 3.

A counterbalancing condition was created by modifying the training materials so that the test items that were grammatical in one language were the ungrammatical test items in the other language (see *Test material*, below). This was accomplished by modifying the non-adjacent dependencies so that in one condition the dependencies were $A_1_B_1$, $A_2_B_2$, $A_3_B_3$, and in the other they were $A_1_B_2$, $A_2_B_3$, $A_3_B_1$ (subscripts indicate the particular item in the word category). This design ensures that any grammaticality effects we find are not due to idiosyncratic properties of the test items.

Test material. As we noted before, there were 27 possible artificial sentences in a particular condition, given that there were 3 different words for each non-dependent position (i.e., positions 1, 2, and 3). To make up the testing material, a total of 54 artificial sentences were used, where 27 sentences were grammatical sentences and the other 27 were ungrammatical for that particular condition. The grammatical sentences in one condition (say, with $A_1_B_1$, $A_2_B_2$, $A_3_B_3$ dependency) was used as ungrammatical sentences in the other condition (with $A_1_B_2$, $A_2_B_3$, $A_3_B_1$ dependency), and vice versa.

Apparatus and Procedure. Participants were tested individually in a quiet testing area. The experiment was conducted on a desktop computer and implemented using PsychToolbox 3 on MATLAB (MathWorks, 2013; Brainard, 1997). Instructions were displayed on the screen and auditory stimuli were presented through headphones. The subject entered responses using the keyboard.

The entire experiment consisted of 3 blocks, each of which consisted of a training phase followed by a test phase. Thus, the 858 English sentences and 432 sentences in the artificial language were evenly divided into 3 parts for the 3 blocks. For testing, 18 sentences (9 grammatical and 9 ungrammatical) were randomly selected from all test sentences (27 grammatical and 27 ungrammatical) for each block. The order of the presentation for these sentences was randomly sequenced for each participant during the testing section. During the training phase, participants listened to the sound stream passively through headphones while the screen was blank. Each block lasted approximately 9 minutes. Immediately after each training block, we showed instructions for the test phase on the screen. The instructions informed the participants that they would hear sound sequences and make judgments about the sequences. There were a total of 18 test trials per block, half of which were grammatical, and the other half ungrammatical. The ordering of the test trials was randomized for each participant.

Participants initiated each test trial. For each trial, we played an artificial language sentence, and asked the participant to indicate whether the sequences were present in the previous section that they had heard. After the sentence was played, a scale was presented and participants were asked to answer the question: “Do you think that you heard this sequence in the previous section?” There were five possible items to choose from, “Definitely”, “Maybe”, “Not Sure”, “Maybe Not”, “Definitely Not”. Participants could click on any of the choices, after which the trial ended and next trial began.

Results

We coded the scale of “Definitely”, “Maybe”, “Not Sure”, “Maybe Not” and “Definitely Not” into numeric values of 4, 3, 2, 1, and 0, respectively. Averaged over the three test blocks, subjects gave numerically higher (i.e., greater endorsement) to grammatical items ($M=2.96$)

compared to ungrammatical test items ($M=2.83$). To ensure that there were no differences across counterbalancing conditions, we fit a 2 X 2 X 3 ANOVA using the R package *ez* (Lawrence, 2015), with counterbalance group as a between-subjects variable and grammaticality and block as within-subjects variables, and subject rating as the dependent measure. There was no effect of counterbalance group or any interactions of group with either other variable, so we collapsed across group for the subsequent analysis. A 2 X 3 (Grammaticality X Block) ANOVA revealed a significant main effect of grammaticality ($F(1,23) = 7.69, p=0.011, \eta_G^2=.069$) and a significant grammaticality by block interaction ($F(1,23) = 4.62, p=0.042, \eta_G^2=.078$), and no main effect of block ($F(1,23) = 0.48, n.s.$). Thus, subjects showed evidence of successfully learning the non-adjacent dependency (Figure 2).

Figure 2. Ratings for all test items (grammatical and ungrammatical) as well as the difference scores for all subjects in Experiment 1. In Figure 2A, each circle represents the mean ratings of a subject for grammatical and ungrammatical items, with a solid line indicating the

mean for each item type. In Figure 2B, each circle represents the difference between ratings (grammatical – ungrammatical) for each subject, with the solid line showing the mean, and shadows showing 95% confidence intervals around the mean. The dotted line is at 0 representing chance.

Discussion

In Experiment 1, we introduced a parsing pattern in English that participants could use to bracket long sequences in the artificial language into four-word long sub-sequences, or sentences. We predicted that the cyclic processing of [*NAME VERB PRONOUN PARTICLE*] in English would cause subjects to apply the entrained pattern to the artificial language, bracketing the longer sequences into subsequences with the structure [*YA_iXB_i*] [*YA_jXB_j*].... That is, the phase of the bracketing was such that the *A_i_B_i* dependency was at the second and fourth position of the same cycle. We further predicted that the bracketing would play a similar role as silences (Gómez, 2002; Peña et al., 2002) in facilitating learning the dependencies. Our results are consistent with these predictions.

By generating the English sentences with a random Markov process, all transitional probabilities between words within sentences and across sentences in the experiment were constant. That is, each English word predicted equally all words in the following category, at least in the context of the experiment. Nevertheless, English speakers automatically apply a structural parse and segment the sequence into sentences, and it is this abstract structure we hypothesize subjects in Experiment 1 applied to the artificial sequences as well.

Unlike in prior studies (Endress & Bonatti, 2007; Newport & Aslin, 2004; Peña et al., 2002; Wang & Mintz, under review), here we found evidence of non-adjacent dependency

learning within continuous sequences of an artificial language. Our claim is that the dependencies were learnable because they occurred within bracketed subsequences which, although abstract, functioned like signal-driven boundaries such as silences.

Given this interpretation, it is then interesting to ask what happens when the top down bracketing information disrupts the NADs, by causing the dependent elements to occur in different subsequences. If our interpretation is correct, learners' ability to detect the dependencies should be hindered. We tested this prediction in Experiment 2.

Experiment 2

Experiment 1 established that learners were capable of learning non-adjacent dependencies when English sentences were interleaved and the dependencies in English were in phase with the dependencies in the artificial language. Experiment 2 also interleaved English sentences with the artificial language, except that the phase of the English sentence and the phase of the artificial dependencies were mismatched. This allowed us to examine whether it was critical in Experiment 1 that the dependencies were contained within the abstract sentence boundaries.

Methods

Participants. Twenty-four undergraduate students at University of Southern California were recruited from Psychology Department subject pool. Half of them participated in each counterbalancing condition. None of the participants participated in Experiment 1.

Training Material. The training materials were similar to those in Experiment 1. The critical difference was at each transition from English to the artificial language: We placed one of the English proper nouns (chosen at random) at the juncture between the last word of an English sentence and the first word of the artificial sentence. Our hypothesis was that, on

hearing the proper noun, subjects would continue processing as they had been, and parse the following three words with the name as part of the entrained four-word sentence structure, resulting in a parse like $[NAME\ X\ A_i\ Y]\ [B_i\ X\ A_j\ Y] \dots$ (see Figure 3). The extra English word results in a shift in the phase of the bracketing of the artificial language sequence. The result of this phase shift was that the A_i element was now in the third position of one cycle, but the B_i element corresponding to it was in the first position of the next cycle. Critically, though, the artificial sequence itself was structured identically to that in Experiment 1. The only difference was in the *English* material, specifically, the extra word at the junctures between English and the artificial language. We hypothesized that this difference would result in subjects inducing abstract boundaries in locations that interrupted the $A_i\text{--}B_i$ dependencies, making them more difficult to learn.

Figure 3. The design of language materials in the training phase of Experiment 2.

Test Materials. The test materials were identical to those in Experiment 1.

Procedure

The procedure was identical to Experiment 1.

Results

Similar to Experiment 1, we coded the scale of “Definitely”, “Maybe”, “Not Sure”, “Maybe Not” and “Definitely Not” into numeric values of 4 through 0. Averaged over the three test blocks, mean ratings for grammatical and ungrammatical test items were 2.94 and 2.86, respectively. To ensure that there were no differences across counterbalancing conditions, we fit a 2 X 2 X 3 ANOVA, with counterbalance group as a between-subjects variable and grammaticality and block as within-subjects variables, and subject rating as the dependent measure. There was no effect of counterbalance group or any interactions of group with either other variable, so we collapsed across group for the subsequent analysis. A 2 X 3 (Grammaticality X Block) ANOVA revealed no significant main effects or interactions (all F s < 1). We therefore conclude that subjects did not learn the non-adjacent dependencies when grammatical entrainment was out of phase with the dependency relations (Figure 4).

Figure 4. Ratings for all test items (grammatical and ungrammatical) as well as the difference scores for all subjects in Experiment 2. In Figure 4A, each circle represents the mean ratings of a subject for grammatical and ungrammatical items, with a solid line indicating the

mean for each item type. In Figure 4B, each circle represents the difference between ratings (grammatical – ungrammatical) for each subject, with the solid line showing the mean, and shadows showing 95% confidence intervals around the mean. The dotted line is at 0 representing chance.

Discussion

In Experiment 2, the phase of the English sentences was mismatched to the phase of the artificial language dependency such that the non-adjacent dependencies crossed the entrained bracketing of the segmented material. In contrast to Experiment 1, where the entrained bracketing was in phase with the dependency structures, participants in Experiment 2 failed to learn the artificial language dependencies.

Our interpretation is that in Experiment 1, subjects' parsing of the the English sentences induced them to impose boundaries in the artificial language at locations that were aligned with the NADs. If English parsing indeed carries over to parsing the artificial language, then the NADs in Experiment 2 occurred in separated subsequences, as in [Kate Y_{A_i}X] [B_iY_{A_j}X]..., and the A_iXB_i dependencies thus were not available to be learned. We thus propose that the failure to learn the dependencies in Experiment 2 was the effect of the entrained bracketing structure segregating the dependent elements across boundaries.

Although we interpret the bracketing effects in Experiments 1 and 2 as operating at the level of the sentence, it is possible that subjects were affected by phrase-level dependencies as well. That is, the artificial language dependency between the second and fourth words was mirrored by the verb+particle dependencies of the phrasal verbs in the English sentences (e.g.,

puts_down). The English dependency is a grammatical one, governed by the hierarchical relationship between the verb and particle, and it could induce subjects to relate the equivalent positions in the artificial language. In addition, even though there are no statistically reliable patterns between the English words in the second and fourth positions within the experiment, subjects' past experience with English plausibly would lead them to represent the verb+particle combinations as familiar patterns, and thus increase their attention to the statistical relationship of the 2nd and 4th positions in the artificial language³. In that case, participants in Experiment 1 may have detected the NAD in the artificial language because of their experience as English speakers with the non-adjacent structure of the phrasal verbs, in addition to the sentence level entrainment. (Note that, even within this alternative, participants would have had to parse the artificial language following the English sentence structure in order to evaluate the analogous positions within a sentence in the artificial language.) To test this alternative, we removed the potential influence of the phrasal verb in Experiment 3.

Experiment 3

To rule out the possibility that the dependencies of the phrasal verbs in Experiment 1 (e.g., *puts_up*) were highlighting the NADs in the artificial language, in Experiment 3 we modified the structure of the English sentences to make the verb+preposition an adjacent sequence (e.g., *John cuts through these*). This modification means that the potentially salient grammatical and statistical dependency of the English phrasal verbs could not lead participants to attend preferentially to the statistical dependencies between the 2nd and 4th words in the

³ We thank an anonymous reviewer for bringing this to our attention.

artificial language. These new sentences were constructed the same way as they were in Experiment 1, where the English sentences did not contain any signal-driven structural cues, and had constant transitional probability between each word within and across sentences. If sentence level entrainment still facilitates learning, we can eliminate the possibility that the facilitation effect in Experiment 1 was a phrase level phenomenon.

Methods

Participants. Twenty-four undergraduate students at University of Pennsylvania were recruited from Psychology Department subject pool. Half of them participated in each counterbalancing condition.

Stimuli and Design. Experiment 3 differed from Experiment 1 only in the design of the English sentences we used as the entrainment materials. Instead of using structures that contain a verb particle dependency between the 2nd and 4th positions of the sentence (e.g., *John puts these down*), we used structures that eliminated such dependencies (e.g., *John cuts through these*).

We recorded speech from the same native English speaker as in the first two experiments, digitized at a rate of 44.1 kHz. Similar to Experiments 1 & 2, we used 2 types of words: English words and novel words. For English words, we used 5 monosyllabic names (*Ross, John, Kate, Nate, Clair*), 5 monosyllabic verbs in 3rd person singular form (*runs, looks, cuts, works, digs*), 5 prepositions (*through, at, for, with, in*), and 5 pronouns (*these, those, this, that, it*). For the non-adjacent dependency, we used the same 12 novel words from the previous experiments.

The words were compressed and spliced as in Experiments 1 and 2.

Training materials. The basic design of the training and testing materials is the same as the materials in Experiment 1, except using the new English words specified in the Stimuli section making up the English sentences (e.g., *John cuts through these*, See Figure 5).

Figure 5. The design of language materials in the training phase of Experiment 3. The English sentence and the artificial dependencies are in phase.

Figure 5 depicts the grammatical bracketing of three English sentences and the entrained bracketing of one artificial sentence in sequence. Unlike Experiment 1 and 2, the verb and its particle no longer occupy the 2nd and 4th positions within a sentence.

A counterbalancing condition was created in the same way as previous experiments.

Test material. The same test material from Experiment 1 was used in Experiment 3.

Apparatus and Procedure. The same apparatus and procedure from Experiment 1 was used in Experiment 3.

Results

We coded the scale of “Definitely”, “Maybe”, “Not Sure”, “Maybe Not” and “Definitely Not” into numeric values of 4 through 0. Averaged over the three test blocks, subjects gave numerically higher (i.e., greater endorsement) to grammatical items ($M=2.94$) compared to ungrammatical test items ($M=2.77$). To ensure that there were no differences across

counterbalancing conditions, we fit a 2 X 2 X 3 ANOVA, with counterbalance group as a between-subjects variable and grammaticality and block as within-subjects variables, and subject rating as the dependent measure. There was no effect of counterbalance group or any interactions of group with either other variable, so we collapsed across group for the subsequent analysis. A 2 X 3 (Grammaticality X Block) ANOVA revealed a significant main effect of grammaticality ($F(1,23) = 5.99, p=0.022, \eta_G^2=.084$), no significant grammaticality by block interaction, and no main effect of block (all $F_s < 1$, n.s.). We conclude that subjects showed evidence of successfully learning the non-adjacent dependency (Figure 6).

To compare Experiments 1 and 3, we conducted mixed-effects linear regressions with data from both experiments. In the regression, the main effects and interactions of experiment and grammaticality were included as fixed effects, and by-subject random intercepts and slopes for grammaticality were random effects. The main effect of grammaticality was significant ($\beta = .148, z=3.73, p<0.001$), but the main effect of experiment was not ($\beta = .045, z=0.30, p=0.763$, ns). Moreover, the interaction was not significant ($\beta = .034, z=0.43, p=0.669$, ns), indicating that there was no evidence that the effects of learning were different between Experiment 1 and 3.

Figure 6. Ratings for all test items (grammatical and ungrammatical) as well as the difference scores for all subjects in Experiment 3. In Figure 6A, each circle represents the mean ratings of a subject for grammatical and ungrammatical items, with a solid line showing the mean for each item type. In Figure 6B, each circle represents the difference between ratings (grammatical – ungrammatical) for each subject, with the solid line showing the mean, and shadows showing 95% confidence intervals around the mean. The dotted line is at 0 representing chance.

Discussion

In Experiment 3, participants successfully learned the non-adjacent dependencies in the artificial language, as they did in Experiment 1. Recall that in Experiment 1, there was the possibility that participants attended to the relationship between the verb and the particle by virtue of their experience of the verb+particle structures in English (e.g., *puts...down*; *turns...in*).

It might have drawn participants' attention to the positions in which NADs occurred in the artificial language. Here we controlled for that possibility, and participants nevertheless learned the dependencies. This suggests that it was the structural information that guided subjects' processing of the artificial language across all the experiments. Note that, although the phrasal verb dependency no longer occupies the 2nd and 4th positions, the verb phrase edges in the English sentences coincide with the dependency in the artificial language. If learners were entrained to this finer grained structure in the English sentences (as in Ding et al, 2015), it could have influenced top-down bracketing of the artificial language sequences. The structural relationship between the verb and the pronoun at the verb phrase edges might even have been made more salient when the sentence final pronoun was an argument, as oppose to an adjunct⁴. These are interesting possibilities which, in our view, are very much in keeping with our broader hypothesis concerning the role of top-down structural information on the detection of NADs. Future research is needed to investigate whether learners are responding to sentence level bracketing only, or are further (or only) guided by more fine-grained structural representations carried over from processing English, such as verb phrase constituents. In either case, the results from Experiment 3 allow us to rule out the possibility that participants' experience with the statistically reliable relationships between the verbs and particles in English was causing them to attend to the statistically reliable NADs in the artificial language. Rather, participants' success in learning the NADs in Experiments 1 and 3, and failure in Experiment 2, was influenced by their processing of the structure of the interleaved English sentences.

⁴ We thank Amy Perfors and an anonymous reviewer for encouraging us to discuss these possibilities.

Having further established the nature of the English structure that learners carried over to processing the artificial language in Experiments 1 and 3, we now address an open question from Experiment 2. In Experiment 2, the occurrence of an English word (a name) at the end of English sentences caused the dependencies in the artificial language to be out of phase with the bracketing carried over from the English sentences. This transition from the English name to nonce words happened unpredictably and not at a sentence boundary. While this unpredictable transition was designed to offset the phase of the nonce words, it may have also induced a violation of the expectation that the name begins an English sentence, and that English words will follow other English words within that sentence. This results in a potential confound, as the lack of NAD learning could be due to general confusion, rather than out of phase sentence bracketing.

To address this in Experiment 4, rather than using English words to manipulate the phase of the artificial language, we used nonce words instead. One of a new class of nonce words was inserted to offset the phase of the upcoming artificial language stream instead of an English name. This way, the transition between the English sentences to the artificial language is at sentence boundaries, while the dependencies to be learned are still out of phase with respect to the English sentences.

Experiment 4

The goal of Experiment 4 was two fold: First, we wanted to have a conceptual replication of learning under the out of phase manipulation in Experiment 2. Second, we wanted to alleviate the concern raised above by showing it is possible to induce the out-of-phase effect while maintaining boundaries between English and novel words at sentence boundaries. We simply inserted a novel word after the end of the English sentences before proceeding with the

artificial language that contains the dependencies. If the prediction from entrainment that the bracketed subsequences determine dependency learning still holds, participants should fail to learn dependencies in Experiment 4 as well.

Methods

Participants. Twenty-four undergraduate students at University of Pennsylvania were recruited from Psychology Department subject pool. Half of them participated in each counterbalancing condition.

Stimuli and Design. Experiment 4 used the same English words and structures as in Experiment 3, and a nearly equivalent artificial language. The only difference was that in the transitions from the English sentences to the artificial language, an extra word (from a new set of three) appeared before the usual artificial language. This replicates the out of phase manipulation in Experiment 2, but removes the potential confound of switching to the artificial language within an English sentence.

Training materials. The method for generating the English sentences was equivalent to Experiment 3. The $Y A_i X B_i$ sequences in the artificial language were also generated as in Experiment 3. However, to offset the phase of the artificial sequences, we placed a randomly-selected word from the new set of novel words Z (*bex*, *nud*, *ghip*) at the juncture between the last word of an English sentence and the first word of the first artificial sentence in the artificial language sequences. Our hypothesis after processing the English sentences, subjects would continue to process four-word ‘sentence’ units in the artificial language, resulting in a parse like $[Z X A_i Y] [B_i X A_j Y] \dots$ (see Figure 7). The extra Z word thus results in a shift in the phase of the bracketing of the artificial language sequence; however, the artificial sequence itself was structured identically to that in Experiment 3. We hypothesized that this difference would result

in subjects inducing abstract boundaries in locations that interrupted the $A_i_B_i$ dependencies, making them more difficult to learn, as in Experiment 2.

Figure 7. The design of language materials in the training phase of Experiment 4. The English sentence and the artificial dependencies are out of phase.

Figure 7 depicts the grammatical bracketing of three English sentences and the entrained bracketing of one artificial sentence in sequence.

A counterbalancing condition was created in the same way as previous experiments.

Test material. The same test material from Experiment 1 was used in Experiment 3.

Apparatus and Procedure. The same apparatus and procedure from Experiment 3 was used in Experiment 4.

Results

We coded the scale of “Definitely”, “Maybe”, “Not Sure”, “Maybe Not” and “Definitely Not” into numeric values of 4 through 0. Averaged over the three test blocks, subjects gave about equal ratings for grammatical items ($M=2.735$) and ungrammatical test items ($M=2.715$). To ensure that there were no differences across counterbalancing conditions, we fit a 2 X 2 X 3 ANOVA, with counterbalance group as a between-subjects variable and grammaticality and block as within-subjects variables, and subject rating as the dependent measure. There was no

effect of counterbalance group or any interactions of group with either other variable, so we collapsed across group for the subsequent analysis. A 2 X 3 (Grammaticality X Block) ANOVA revealed no significant main effect of grammaticality or block (F 's < 1), and no significant grammaticality by block interaction ($F(1,23)=2.58$, $p=0.12$, *ns*). We conclude that subjects showed no evidence of learning the non-adjacent dependency (Figure 8).

Figure 8. Ratings for all test items (grammatical and ungrammatical) as well as the difference scores for all subjects in Experiment 4. In Figure 8A, each circle represents the mean ratings of a subject for grammatical and ungrammatical items, with a solid line indicating the mean for each item type. In Figure 8B, each circle represents the difference between ratings (grammatical – ungrammatical) for each subject, with the solid line showing the mean, and shadows showing 95% confidence intervals around the mean. The dotted line is at 0 representing chance.

Discussion

We conclude that using a novel word in Experiment 4 to offset the phase disrupts learning, similar to Experiment 2. Using an English name or using a novel word achieved the same effect. This finding rules out the possibility that the failure to learn in Experiment 2 is due to the fact that the transition from English words to nonce words was not at a sentence boundary in Experiment 2.

More generally, learning failed again when the phase of the English sentences was mismatched to the phase of the artificial language dependency in Experiment 4, similar to the results in Experiment 2. These results provide further evidence in favor of our claim that subjects applied an abstract sentence bracketing on the artificial language, and in the case of out of phase alignment, this bracketing disrupts learning. The out of phase alignment resulted in subjects parsing the artificial language into subsequences like $[ZYA_iX]$ $[B_iYA_jX]$..., which apparently disrupts subjects ability to learn the A_iXB_i dependency, which is now interrupted by an induced boundary. Thus, in Experiments 2 and 4, in contrast with Experiments 1 and 3, the entrained top-down structures are out of alignment with the statistical patterns in the artificial language, and this interferes with subjects' ability to learn the statistical patterns, even though the actual bottom-up statistical patterns are identical across all the experiments.

General Discussion

In this paper we report the first findings we are aware of that show that adult learners can learn dependencies between non-adjacent words in a continuous stream of an artificial language. Moreover, we showed that learning was influenced by abstract bracketing information that, in our experiments, was provided through entrainment to grammatical structures of English

passages. As we mentioned earlier, results from prior studies (Newport & Aslin, 2004; Peña et al., 2002) suggested that having silences surrounding the dependencies was critical to the learning of (syllable-level) non-adjacent patterns. Our findings demonstrate that silences are not necessary, and they suggest that any source of bracketing information can facilitate the detection of dependency patterns, so long as the bracketing contains the dependencies, rather than breaking them up and segregating the dependent elements from each other. Our findings thus suggest that in prior experiments, silences facilitated non-adjacent dependency learning because they provided a smaller constituent that contained the critical elements. These findings have a number of implications for the interpretation of prior research. They also provide important new constraints on theories about the nature of human statistical learning. We now discuss these consequences in turn.

The Role of Edges In Learning Non-Adjacent Dependencies

Some existing theories posit that perceptual or memory primitives guide aspects of statistical learning, and more specifically, that edge-based computations are critically required for learning non-adjacent dependencies (Endress & Bonatti, 2007; Endress et al., 2009; Peña et al., 2002). In the studies that motivated that view, and in studies of non-adjacent dependency learning more broadly (Gómez, 2002; Gómez & Maye, 2005; Romberg & Saffran, 2013), the edges were “signal-driven” (Peña et al., 2002) silences surrounding the dependencies. Our findings bear on this theory in two important ways. First, as we just discussed, we show that silences or signal-driven information *per se* are not critical, and non-adjacent dependency learning is facilitated by top-down bracketing cues as well. On the surface, this outcome is consistent with the general idea that edges are important for detecting dependencies (Endress & Bonatti, 2007; Endress et al., 2009; Peña et al., 2002), but takes a broader view on what

constitutes an edge. However, it is also consistent with a more general account whereby edges (however derived) focus statistical learning mechanisms on the locations in the longer sequence that contain the dependencies. On this account, special rule-learning mechanisms (specific to edge-aligned material) are not required to learn the non-adjacent dependencies; rather, adjacent and non-adjacent patterns may be learned by similar mechanisms, but the mechanisms operate within constituency boundaries that are marked by edges, whether derived from bottom-up signal-driven information, or top-down, abstract structural information.

The Importance Of Constituency

A corollary to this view concerning edges is that lower-level statistical learning is constrained by constituent structure, at least with respect to quasi-linguistic material. The artificial languages across the four experiments were identical in the linear sequences of words (down to the level of the auditory signal), but they differed in their abstract constituent structure, which arose as a consequence of processing the interspersed English sentences. When the dependencies were contained in constituents (either a verb phrase or a sentence), subjects learned them (Experiments 1 and 3); when the dependencies spanned constituent boundaries, there was no evidence of learning (Experiments 2 and 4). Regardless of whether the entrainment effect comes from the sentential or phrasal level, the important finding here is that statistical learning was constrained by a top-down structure building process. Of significance is the bracketing of artificial languages from the repeated exposure to English materials with the same structure. This structure building process determined what statistics learners had access to (i.e. that they had access to statistics within subsequences and no access to statistics across subsequences), as identical surface-level statistical information was available to learners in both sets of experiments.

The idea that constrained statistical/distributional analysis could be involved in grammatical acquisition has an important history in language acquisition research. For example, Pinker (1984, pp. 40-42) proposed that children could acquire aspects of grammar using distributional analyses that are constrained to operate within syntactic constituents. Kimball (1973) noted that the high frequency function words that occur at phrase boundaries make them salient cues to constituent boundaries. Early experiments with artificial grammar learning in adults demonstrated benefits to grammatical acquisition when high frequency words or other phrase boundary cues were provided in the input (Morgan, Meier, & Newport, 1987; Valian & Coulson, 1988). The experiments reported here are consistent with those findings and provide further evidence that human learners are highly responsive to constituency cues. Our study goes further, however, in showing that learners use boundary cues to constrain their statistical analyses even when they are abstract, top-down structural cues that result from processing distal information (here, the English sentences).

Finally, we have focused on the sentence level information as the source of entrainment and constituency. While it is plausible that phrase level information could have effects as well, we did not find support for this in Experiment 3. Without the presence of the phrasal level non-adjacent dependency in the English sentences (e.g., *picks_up*), participants in Experiment 3 were still able to learn the non-adjacent dependency in the artificial language. We also found no difference in learning between Experiment 1 and 3, supporting the claim that the phrase level information in Experiment 1 did not have an effect of facilitating non-adjacent dependency learning.

Concluding Remarks

Some readers might object that various aspects of our design are artificial, especially in terms of the temporal nature of the English and artificial language material. For instance, natural language is not nearly so regimented with respect to cyclic rhythmic structures. However, our goal was to use the rhythmic properties as a tool to stimulate syntactic transfer from processing of known structures to novel ones. By employing this entrainment method that has been shown to evoke abstract structural processing (Ding et al., 2015), we were able to stimulate processing mechanisms in subjects that led them to bracket the artificial language using top-down information, rather than signal-driven information in the form of silences or other prosodic cues, and thereby learn the non-adjacent patterns. This allowed us to draw broader conclusions about the role silences played in prior studies, and the facilitatory role of boundaries and constituency more broadly. Specifically, we claim that silences provided bottom-up bracketing information that simplified the process of detecting dependencies in longer sequences, by containing the dependency within a structured sequence, just as the entrained sentence structure did in Experiment 1, but not in Experiment 2. We know of no prior research that has shown that structural entrainment affects learning. This raises the intriguing possibility that similar processes are engaged in natural language learning. For example, structural entrainment can be viewed either as prediction signals for segmentation, or more broadly as a kind of syntactic priming (Pickering & Branigan, 1998), where prior structures organize the interpretation of novel material.

In conclusion, our study is the first that we know of to show that human adults can learn non-adjacent lexical patterns in a novel language when the critical sequences are not pre-segmented by silence. We showed that learners use abstract top-down structural information to bracket continuous sequences of words into smaller constituents, and that when dependencies are

contained within those constituents, they are learned. Our findings therefore do not support theories that require the non-adjacent dependent elements to occur at edges defined by silences. They also argue against theories that associate special learning mechanisms with edge positions, at least with respect to non-adjacent dependency learning. We also show, contrary to some accounts (e.g., Gómez, 2002) that non-adjacent dependencies can be learned even when the intervening position shows limited variability. We propose that the benefit prior research has shown for silences derived from a bracketing of the input that focuses statistical learning mechanisms on constituents that contain the dependency. We propose that thinking about the role of bracketing is a useful way of gaining insights into the learning of linguistic dependencies, and human statistical learning more broadly.

Appendix: Example Sequences from Experiments

The examples below are to show some of the possible sequences in the experiments. The novel words making up the dependency are italicized for demonstration purposes only: none of the auditory stimuli of the words in the experiments are marked in any way.

Experiment 1: Brian turns these down John puts that on Nate has it up Clair lets this in Kate keeps those on *pel swech dap ghire tink voy tood jub blit rud wesh tiv* Nate puts those on Kate turns these off ...

Experiment 2: Brian turns these down John puts that on Nate has it up Clair lets this in Kate *pel swech dap ghire tink voy tood jub blit rud wesh tiv* keeps those on Nate puts those on Kate turns these off ...

Experiment 3: Ross runs at these John looks through that Nate cuts through those Clair digs at this Kate works with those *pel swech dap ghire tink voy tood jub blit rud wesh tiv* Nate cuts through those Kate works in these...

Experiment 4: Ross runs at these John looks through that Nate cuts through those Clair digs at this Kate works with those *bex pel swech dap ghire tink voy tood jub blit rud wesh tiv* cuts through those Kate works in these...

References

- Boersma, Paul (2001). Praat, a system for doing phonetics by computer. *Glott International* (2001): 341-345.
- Ding, N., Melloni, L., Zhang, H., Tian, X., & Poeppel, D. (2015). Cortical tracking of hierarchical linguistic structures in connected speech. *Nature Neuroscience*, 19, 158-64. doi:10.1038/nn.4186
- Endress, A. D., & Bonatti, L. L. (2007). Rapid learning of syllable classes from a perceptually continuous speech stream. *Cognition*, 105, 247-299.
- Endress, A. D., Nespore, M., & Mehler, J. (2009). Perceptual and memory constraints on language acquisition. *Trends in Cognitive Sciences*, 13, 348-353.
- 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.
- Gómez, R., & Maye, J. (2005). The developmental trajectory of nonadjacent dependency learning. *Infancy*, 7, 183-206.
- Kimball, J. (1973). Seven principles of surface structure parsing in natural language. *Cognition*, 2, 15-47.
- Lawrence, M. A. (2015). ez: Easy Analysis and Visualization of Factorial Experiments. R package version 4.3. <https://CRAN.R-project.org/package=ez>.
- Luo, H., & Poeppel, D. (2007). Phase patterns of neuronal responses reliably discriminate speech in human auditory cortex. *Neuron*, 54, 1001-1010.

- Mintz, T.H. (2003). Frequent frames as a cue for grammatical categories in child directed speech. *Cognition*, 90, 91-117.
- Mintz, T. H., Wang, F. H., & Li, J. (2014). Word categorization from distributional information: Frames confer more than the sum of their (Bigram) parts. *Cognitive Psychology*, 75, 1-27.
- Morgan, J. L., Meier, R. P., & Newport, E. L. (1987). Structural packaging in the input to language learning: contributions of prosodic and morphological marking of phrases to the acquisition of language. *Cognitive Psychology*, 19, 498–550.
- Newport, E.L., & Aslin, R.N. (2004). Learning at a distance: I. Statistical learning of non-adjacent dependencies. *Cognitive Psychology*, 48, 127-162.
- Pelucchi, B., Hay, J. F., & Saffran, J. R. (2009). Statistical learning in a natural language by 8-month-old infants. *Child Development*, 80, 674–685.
- Peña, M., Bonatti, L. L., Nespor, M., & Mehler, J. (2002). Signal-driven computations in speech processing. *Science*, 298, 604-607.
- Pickering, M. J., & Branigan, H. P. (1998). The representation of verbs: Evidence from syntactic priming in language production. *Journal of Memory and Language*, 39, 633-651.
- Pinker, S. (1984). *Language Learnability and Language Development*. Cambridge, MA: Harvard University Press.
- Romberg, A. R., & Saffran, J. R. (2013). All together now: Concurrent learning of multiple structures in an artificial language. *Cognitive Science*, 37, 1290-1320.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274, 1926-1928.

- Santelmann, L. M., & Jusczyk, P. W. (1998). Sensitivity to discontinuous dependencies in language learners: Evidence for limitations in processing space. *Cognition*, 69, 105-134.
- Slobin, D. I. (1973). Cognitive prerequisites for grammar. In C. Ferguson & D. I. Slobin, eds., *Studies in child language development*. New York: Holt, Rinehart and Winston.
- Thompson, S.P., & Newport, E.L. (2007). Statistical learning of syntax: The role of transitional probability. *Language Learning and Development*, 3, 1-42.
- Valian, V. V., & Coulson, S. (1988). Anchor points in language learning: The role of marker frequency. *Journal of Memory and Language*, 27, 71–86.
- Wang, F.H., & Mintz, T.M. (in press). Learning Non-Adjacent Dependencies Embedded in Sentences of an Artificial Language: When Learning Breaks Down. *Journal of Experimental Psychology: Learning, Memory, and Cognition*
- Wang, Y., Ding, N., Ahmar, N., Xiang, J., Poeppel, D., & Simon, J. Z. (2012). Sensitivity to temporal modulation rate and spectral bandwidth in the human auditory system: MEG evidence. *Journal of Neurophysiology*, 107, 2033-2041.