Learning Non-Adjacent Dependencies Embedded in Sentences of an Artificial Language:

When Learning Breaks Down

Felix Hao Wang and Toben H. Mintz

University of Southern California

Author Note

Felix Hao Wang, Department of Psychology, 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 37th annual meeting of the Cognitive Science Society in Pasadena, CA.

Abstract

The structure of natural languages give rise to many dependencies in the linear sequences of words, and within words themselves. Detecting these dependencies is arguably critical for young children in learning the underlying structure of their language. There is considerable evidence that human adults and infants are sensitive to the statistical properties of sequentially adjacent items. However, the conditions under which learners detect non-adjacent dependencies (NADs) appears to be much more limited. This has resulted in proposals that the kinds of learning mechanisms learners deploy in processing adjacent dependencies are fundamentally different from those deployed in learning NADs. Here we challenge this view. In four experiments, we show that learning both kinds of dependencies is hindered in conditions when they are embedded in longer sequences of words, and facilitated when they are isolated by silences. We argue that the findings from the present study and prior research is consistent with a theory that similar mechanisms are deployed for adjacent and non-adjacent dependency learning, but that NAD learning is simply computationally more complex. Hence, in some situations NAD learning is only successful when constraining information is provided, but critically, that additional information benefits adjacent dependency learning in similar ways.

Keywords: Artificial language learning; Adjacent dependencies; Non-adjacent dependencies; Grammatical acquisition; Sequence Learning

Learning Non-Adjacent Dependencies Embedded in Sentences of an Artificial Language:
When Learning Breaks Down

Sentences in natural languages contain grammatical dependencies, such as those that arise from agreement marking between the sentence subject and the verb. Sometimes these dependencies hold between adjacent words (or morphemes), and sometimes the dependencies are non-adjacent. For example, the dependency between the singular subject *child* and the agreeing inflected verb *runs* in *the child runs* is an adjacent dependency with respect to lexical linear position, whereas it is non-adjacent in, *the child always runs*.

There has been considerable interest in investigating learning mechanisms that could detect these dependencies in linear sequences within spoken utterances. Since sequential dependency often result from underlying syntactic hierarchical relationships, mechanisms that detect surface dependencies could be useful to children in acquiring their language's syntax. They could also aid proficient language users in building syntactic parses. There is now considerable evidence of cross-species sensitivity to adjacent co-occurrence dependencies in sequential material, including human infants (Saffran, Newport & Aslin, 1996; Aslin, Saffran & Newport, 1998), human adults (Romberg & Saffran, 2013), monkeys (Hauser, Newport & Aslin, 2001), and rats (Toro & Trobalón, 2005), among others. On the other hand, experiments into sensitivities of non-adjacent dependencies (NADs) have yielded more variable results, and suggest that conditions under which NADs can be learned is more constrained. For instance, some studies suggest that there needs to be considerable variability in the items that occur in the position between the dependent items (Gómez, 2002; Gómez & Maye, 2005), that dependent items need to be bounded by silence (Newport & Aslin, 2004; Peña, Bonatti, Nespor, & Mehler, 2002), or that the dependent items need to share other linguistic features (Newport & Aslin,

2004). Moreover, the mechanism for learning adjacent and non-adjacent dependencies, especially the relationship between the learning mechanisms, is still unclear. Some researchers have characterized learning adjacent dependency as a relatively unconstrained implicit statistical learning process that computes co-occurrence statistics over all adjacent syllables (e.g., Aslin et al., 1998; Saffran et al., 1996). On the other hand, the mixed success on non-adjacent dependency learning has led to proposals that there are additional learning constraints for learning non-adjacent dependencies (Peña et al., 2002; Gómez, 2002; Santelmann & Jusczyk, 1998). Some proposals further claim that NAD learning is a form of rule learning, and fundamentally different from the bottom-up statistical learning involved in learning adjacent dependencies (Endress & Bonatti, 2007; Endress, Nespors, & Mehler, 2009; Peña et al., 2002). In this paper, we present empirical evidence that adjacent and non-adjacent dependency learning may involve learning mechanisms that are more similar than some previous proposals claim, and that their apparent difference may arise simply from greater resource demands in acquiring NADs. Before presenting our experiments, we review in a bit more detail the theoretical proposals regarding adjacent and non-adjacent dependency learning.

Adjacent Dependency Learning

Most accounts of adjacent dependency learning propose that organisms track the co-occurrences of immediately adjacent items—syllables, in most artificial language materials—and calculate statistics based on these co-occurrence events. Saffran et al. (1996) and others discussed their experiments and findings in terms of *transitional probability* (TP) statistics: The probability that the occurrence of a syllable A will be immediately followed by syllable B is the *forward* TP of the sequence AB. Other researchers have discussed both forward and *backward* TPs—the probability that B was immediately preceded by A in the previous example (e.g.,

Pelucchi, Hay, & Saffran, 2009a; Romberg & Saffran, 2013). Other researchers have discussed *mutual information* as the possible statistic (Swingley, 2005), which also considers forward and backward prediction strength. Regardless of the particular statistic being considered, any two adjacent items are judged to be dependent following the magnitude of the statistic.

In practice, many experiments involving adjacent dependencies have explored this kind of statistical learning as a method for segmenting speech (Aslin et al., 1998; Saffran et al., 1996; Pelucchi, Hay, & Saffran, 2009b). Consequently, in those experiments, the syllable streams consisted of long, continuous sequences; positional information—e.g., sequence-initial, sequence-final, etc.—was not a stimulus dimension that was considered to be relevant. Sequential statistics were calculated for all adjacent items, regardless of their position in the sequence. As we will discuss further, this point is important in comparing adjacent to non-adjacent dependency learning.

Non-Adjacent Dependency Learning

Learners could, in principle, deploy the same mechanisms to learn NADs. However, one would expect that detecting and encoding NADs to be slower and more error-prone compared to detecting and encoding adjacent dependencies. This is because the possible dependent elements need to be kept in memory over longer distances (Baddeley, Thomson & Buchanan, 1975). An additional complicating factor has to do with the *span* of the NAD. If a given item could be in a non-adjacent dependency relationship with an item down-stream, it need not be restricted to the item two elements down (i.e., with just one intervening element). Thus, formally, the number of possible dependencies that are rooted in one initial element, a_1 , is bounded only by the length of the sequence. Moreover, the next element, a_2 , is itself a potential first element in an NAD, and therefore downstream relationships with a_2 would need to be tracked along with those of a_1 , and

so on. Thus, without some external or internal constraints on the spans that a learner considers, the statistical learning mechanism just discussed would be limited in detecting NADs by constraints on memory and processing. Of course, those resource constraints could be the manifestation of internal constraints on NAD spans. In any case, there is clearly greater complexity involved for strictly bottom-up statistical mechanisms when the domain of the dependency is non-adjacent, compared to immediately adjacent domains.

Consistent with the expectation that statistical learning of NADs is substantially more difficult, Newport & Aslin (2004, Experiment 1) found that learners failed to learn NADs between syllables in an uninterrupted stream of syllables¹. But other theories of NAD learning have been proposed that posit different mechanisms from the variety typically proposed for adjacent dependency learning (that is, ones based on sequential co-occurrence statistics). One prominent theory, which we will call the *edge-based computation* theory of non-adjacent dependency learning, argues that NADs must be bounded by silences or any kinds of “edges” to be learnable (Peña et al., 2002). When silences occur at the boundaries of NADs in an otherwise continuous sequence, they create sub-sequences, the edges of which become salient to learners. When the dependent elements are at the edges, their salient position in the sequence makes dependencies between them easier to learn, compared to when they are no more salient than surrounding elements in a continuous sequence of syllables (Peña et al., 2002). Endress et al. (2009) further proposed that edge based computation is crucial for NAD learning because edges provide start and end markers that are computationally privileged. Their privileged status results

¹ Newport and Aslin (2004) reported that there is non-adjacent dependency learning between non-adjacent segments in adjacent syllables, when the segments matched in type (i.e., vowel and vowel, consonant and consonant). Newport and Aslin explained their findings in the framework of autosegmental phonology, in which consonants and vowels reside on separate representational tiers; on those tiers, the dependencies subjects learned where, in fact, adjacent.

in the rapid acquisition of dependency rules using mechanisms that are different from those that involve the accumulation of statistical evidence (Endress & Bonatti, 2007).

The studies motivating the proposals just discussed involved learning dependencies between syllables in continuous sequences, and were operationalized as a task of segmenting words from these sequences, as well as learning regularities—the NADs—within the words. Other studies implemented similar designs, but were operationalized as detecting dependencies between words in three-word utterances in an artificial language. Notably, in those studies there were relatively long pauses between each three-word utterance (and shorter silences between words within utterances), placing the dependent words at the boundaries defined by the longer silences (Gómez, 2002; Gómez & Maye, 2005; Gómez, Bootzin & Nadel 2006; Romberg & Saffran, 2013). Most of the experiments in those studies reported successful NAD learning (although some were subject to additional constraints on the variability of the items that occurred in the intervening position of a dependency, e.g., Gómez, 2002).

In sum, the range of studies just discussed suggests that acquiring NADs is conditioned on whether the trigrams containing the dependency are pre-segmented by boundaries of silence. These findings are consistent with edge-based accounts of non-adjacent dependency learning, contrasting with the classical finding that adjacent dependencies can be learned based purely on distributional analysis, without resorting to any type of external cue (Saffran et al, 1996). On the other hand, these findings are also consistent with an explanation related to our earlier discussion of the increased computational complexity of non-adjacent compared to adjacent dependency learning. Recall that the increase in memory and processing resources was largely linked to the increase in the possible spans of NADs. It could simply be that the boundaries/edges signaled by silences provides an explicit bound on the possible spans, and thus makes it easier for learners to

reliably detect and encode the NADs. Under this account, bottom-up statistical mechanisms could account for adjacent and non-adjacent dependency learning, and there would be no need to posit special edge-based mechanisms for acquiring NADs.

If similar statistical learning mechanisms are involved in learning adjacent and non-adjacent dependencies, then it follows that silences that provide boundaries in the signal might aid adjacent dependency learning as well. As far as we know, no study has assessed the benefit of signal-level boundaries in learning adjacent dependencies—perhaps because the learning is generally very robust—but if we can find a situation in which adjacent dependencies are hard to learn, then we might see an advantage when the signal provides boundary cues. Such a finding would lend further support to the idea that similar mechanisms are involved in all types of sequential dependency learning.

Experimental Overview

In this paper, we report on such a situation. In the study we report here, we tested word-level adjacent and non-adjacent dependency learning, manipulating whether or not the sequences in which the dependencies occurred were embedded within longer utterances in an artificial language. In the experiments with non-embedded dependencies (2 & 4), we replicated the stimulus presentation properties of prior studies (e.g., Gómez, 2002; Romberg & Saffran, 2013), and therefore bounded the dependencies within longer silences. In the experiments with embedded dependencies (1 & 3), we placed the dependencies in sequences of unrelated, “buffer” words. One motivation for embedding the NADs in longer sequences was to test whether the difficulties in detecting NADs in syllable-level tasks was also apparent when it came to detecting dependencies between words. This is important information for understanding how learners could detect dependencies in natural language that do not occur at silence boundaries, which

might constitute a significant majority of cases. To preview, we found that when NADs were embedded in longer utterances, subjects failed to learn them (Experiment 1). However, replicating prior studies, when the same NADs were bounded by longer silences, subjects could learn them (Experiment 2). This allowed us to compare the effects of embedding on adjacent dependency learning. We found that when word-level adjacent dependencies were embedded in longer sequences, subjects did not learn them (Experiment 3), but when they were bounded by silence, subjects learned them (Experiment 4)

Experiment 1

The goal of Experiment 1 was to assess word-level non-adjacent dependency learning, when the dependencies were embedded in longer sequences, and where the dependent elements did not occur at sequence edges. We based the design of the material on experiments in which non-adjacent dependencies were learnable by adults and infants (Gómez, 2002). Specifically, the non-adjacent dependencies occurred in trigrams, schematized as A_iXB_i (where the letters stand for classes of words, and the same A and B words always co-occur) adhering to the same statistical structure that lead to learning in Gómez (2002). However, we embedded the trigrams in longer word sequences in the current experiment. If non-adjacent dependencies need to be located at sequences edges to be detectable, learners would not be predicted to learn them in these embedded contexts.

In addition to testing the effects of embedding on learning non-adjacent dependencies, we tested a hypothesis from Gómez (2002). Gómez proposed that learners only notice the A_B dependencies when the adjacent AX and XB sequences are too variable to track. The proposal was based on findings that the success in learning the non-adjacent patterns was affected by the size of the X category, with poor performance when X was restricted to a set of 2 words.

However, Gómez did not directly assess how well subjects actually tracked adjacent dependencies, so the proposed tradeoff between adjacent and non-adjacent learning was not directly tested. More recent data (Romberg & Saffran, 2013) suggest that both kinds of dependencies are learnable together (although perhaps by different mechanisms), suggesting that learning one kind of dependency does not occur at the cost of learning the other. But in the Romberg and Saffran study, as in Gómez (2002), subjects heard pre-segmented trigrams. While we had no reason to believe that the relationship between adjacent and non-adjacent dependency learning should change in embedded context, since this is a relatively unexplored area, we collected data on what subjects learned about the adjacent patterns as well. To maximize the chance of successful learning of non-adjacent dependencies, and to provide a strong test of the hypothesis that non-adjacent and adjacent dependency learning are inversely related, we adopted the version of Gómez's design in which the middle elements alternated between 24 different words, increasing the likelihood for detecting non-adjacent dependencies and making the adjacent patterns highly variable.

Methods

Participants. This study was approved by the USC University Park Institutional Review Board. Twenty-six undergraduate students at University of Southern California recruited from psychology subject pool participated. Subjects were divided equally into two counterbalancing conditions that determined the training language.

Stimulus creation and design. The stimuli were recorded by a female American English speaker for a different experiment (Mintz, Wang, & Li, 2014). The speaker pronounced one word at a time in list citation prosody, and words were digitized at 44.1 kHz for later processing. We then digitally spliced the recording into individual word files that began at the onset of each

word. Word files generated were all shorter than 0.8 s, so silences were added such that concatenating word files results in words occurring every 0.8 seconds within a “sentence”.

Between each artificial sentence, there was an additional 0.8 seconds of silence. The words used in the experiments and the class we assigned them to can be found in Table 1.

The non-adjacent dependencies in the artificial language followed the stimulus design in Gómez (2002). There were three A_iB_i frames and 24 words in the intervening position (i.e., $A_iX_jB_i$, where $i=1-3$ and $j=1-24$). The two counterbalancing languages were created by taking three pairs of A_B for the three frames in one training language (A_1B_1 , A_2B_2 , A_3B_3), and three different pairs for the other training language (A_1B_2 , A_2B_3 , A_3B_1). This design allowed the use of the same set of test items for the two counterbalancing training languages, such that grammatical dependencies for one language were ungrammatical for the other.

 ----- Table 1 Here -----

To create the materials for the learning phase, each frame was presented 6 times with each X word, for a total of 144 presentations of each frame and 18 presentations of each X word. In addition, we added 1-3 *buffer words* both preceding and following each AXB trigram. We call each sequence of (initial buffer word(s) + A + X + B + final buffer word(s)) a sentence. Again, within a sentence a word occurred every 0.8 seconds, with an additional 0.8 second of silence between sentences. Buffer words consisted of 16 words that were not used in either frames position (A , B words) or the intermediate position (X words). For each sentence, a random shuffle of these words was generated, and two sets of 1-3 words were selected from the list be added before and after the non-adjacent dependency. The cardinality of the set was determined

randomly at each draw. This method ensured that no words repeated within a sentence. Further, the AXB trigram could occur only after at least 1 word was presented, and the end of the trigram (B word) would not be the last word within a sentence to be heard. Since the buffer words were randomly distributed, they provided no distributional information about any other part of the sentence. This design is illustrated in Table 2.

 ----- Table 2 Here -----

Apparatus. All the experiments in this paper were conducted in individually isolated sound-attenuating booths, and all participants instructed to wear headphones throughout the experiment. The experiment was implemented in Matlab 2013B with Psychtoolbox 3.

Procedure. The experiment was composed of two phases: learning and testing. The participants were asked to listen to the material to “learn a language”, and they were told that they would answer questions about the language after hearing it.

In the learning phase, participants heard 9 blocks of 43 sentences and 1 block of 45 sentences, selected from the complete set of 432. Each block lasted four and a half minutes, with the entire learning phase being around 50 minutes. When each block ended, a quiz question appeared on the screen: “What was the last word that you heard?” A numbered list of words was displayed, and participants were prompted to press a number key to respond. After the subject answered, the screen went blank and the auditory stimuli resumed. These ‘quizzes’ were designed to encourage subjects to attend to the material.

In the testing phase, we tested adjacent order knowledge before testing non-adjacent dependencies. Knowledge of adjacent order was tested by presenting bigrams that were part of

an AXB sequence (e.g., AX, or XB), henceforth *grammatical*. Foil items were made up by presenting the reverse order of the bigrams, for example, XA or BX, henceforth *ungrammatical*. Note that this is an atypical way of testing adjacent dependencies, not only because the elements making up the ungrammatical items did not appear in the training materials (whereas grammatical items did), but also that the order of classes of those elements in the ungrammatical items was also violated (e.g., XA and BX words never occurred).

There were a total of 12 bigram test items, 6 grammatical and 6 ungrammatical. There were equal numbers of test items with A as there were with B, and there were equal number of grammatical and ungrammatical items with A or B, which means that we tested 3 AX and 3 XB pairs (both grammatical), and 3 XA and 3 BX pairs (both ungrammatical). In order to not induce test effects, the same AX was not tested in reverse (for example, if A_1X_5 was tested, X_5A_1 was not). These properties held equally for both counterbalancing conditions. Whether a particular X word was in a grammatical or ungrammatical test item was counterbalanced across subjects.

There were 6 non-adjacent dependency test items, 3 grammatical and 3 ungrammatical. The grammatical items were taken from the items presented in the learning phase (e.g., $A_1X_1B_1$, $A_2X_1B_2$, $A_3X_1B_3$), while the ungrammatical foils were taken from the non-exposed language (e.g., $A_1X_1B_2$, $A_2X_1B_3$, $A_3X_1B_1$). Thus, the two adjacent bigrams in each of the grammatical and ungrammatical foils were equally familiar, as they all occurred with equal frequency in the familiarized language. The only difference between the two test types was whether the non-adjacent dependency was one that occurred in the familiarized language.

Note that the adjacent and non-adjacent test items are structured differently. The non-adjacent items directly test for knowledge of a dependency (i.e., co-occurrence) between the A and B items (that A_i goes with B_i). In contrast, the adjacent order knowledge only requires

subjects to provide judgments about the relative order of two items that co-occurred (i.e., were dependent) in the training material. We tested subjects' knowledge of adjacent order first, since the test items in the trials testing adjacent dependencies provide no information that could influence subjects' knowledge of the non-adjacent dependencies. However, testing the non-adjacent dependencies first would further expose subjects to the adjacent patterns they heard during the training phase, since all AX and XB subsequences in the non-adjacent test items conform to the trained adjacent patterns.

Subjects heard testing sequences through headphones and on each trial answered the question on the screen: "Did you hear this sequence before?" A Y/N keyboard response was collected for this question. After a short pause, the next test trial began. When the study ended, we thanked our participants and debriefed them about the purpose of the study.

Results

In Experiments 1 through 4, all quizzes were answered at above chance levels, so they will not be mentioned again.

For test items, we analyzed trials testing non-adjacent dependencies and adjacent order information separately. We conducted mixed effects logistic regressions with the binary response (yes/no) as the dependent measure, and counterbalancing group as a between-subjects fixed effect. The results we report here included by-subject intercepts as a random effect; we also ran models with random item intercepts, but since they did not improve model fit, we do not report them here. (Note that random slopes are not meaningful for subject level random effects, since counterbalancing group is a between-subjects variable.)²

First, we report learning performance of the adjacent order information. Again, the

² The Stata syntax for this equation: `xtmelogit outcome i.group || subject:`

purpose of testing knowledge of adjacent items was to rule out the possibility that learning of adjacent dependencies hinders learning of non-adjacent dependencies (Gomez, 2002). Our results showed that there was no knowledge for adjacent items for Experiment 1. Specifically, the mean accuracy for the test trials for adjacent order tests was at 52.89% (Figure 1). The two counterbalancing groups were not statistically different ($\beta=0.044$, $z=0.19$, $p=0.847$; n.s.); the intercept was also not significant ($\beta=0.095$, $z=0.62$, $p=0.537$, n.s.), suggesting that the performance was at chance. Given that A words occur before X words and B words after X words has occurred the same number of times as the non-adjacent dependency, this data thus show that participants did not learn even basic sequential order information about the adjacent dependencies.

----- Figure 1 Here -----

Turning to non-adjacent dependencies, we found no evidence that participants learned embedded non-adjacent dependencies. The mean accuracy for the test trials for non-adjacent dependencies was at 52.57%. The two counterbalancing groups were not statistically different ($\beta=-0.088$, $z=-0.27$, $p=0.786$; n.s.); the intercept was also not significant ($\beta=0.143$, $z=0.65$, $p=0.513$, n.s.), suggesting that the performance was at chance. Mean accuracy of performance for non-adjacent dependency learning by experiment and confidence intervals are shown in Figure 2.

----- Figure 2 Here -----

Of course, a failure to find non-adjacent dependency learning is not, by itself, positive evidence that subjects did not learn the non-adjacent dependencies. In order to test the hypothesis that there was no learning, we used a Bayesian statistical analysis technique to compute the odds on the null (Gallistel, 2009). Briefly, this can be achieved by specifying both a null and alternative prior distribution on accuracy, and computing posterior probabilities for each hypothesis by computing the overall likelihood of the prior given the data. The ratio of these posterior probabilities is the odds ratio between the null posterior and the alternative posterior. If this odds ratio is bigger than 1, we find evidence for the null. If the odds ratio is smaller than 1, we find evidence against the null.

We first used this method pooling data from all the participants in the experiment. Figure 3 demonstrates the change of the Bayes Factor as the prior probability range from at chance to different from chance. In our data, the Bayes Factor for the null is 30.86 for the adjacent order knowledge test items, and 18.41 for the NAD test items, strongly suggesting that learning did not occur. Furthermore, we applied the same method to data from each participant, examining whether all participants showed the same pattern (Figures 4C & 4D). For each subject, the null was pitted against two different alternatives, following Gallistel (2009): The prior probability anywhere between 0.5 and 1 (above chance in the general sense), and the prior probability anywhere between 0.5 and 0.65 (slightly above chance). Most bars show positive values, indicating that most participants showed evidence that they did not learn the non-adjacent dependency or the adjacent order. We concluded that the null result was not due to pooling from a population with a lot of individual differences.

----- Figure 3 Here -----

Discussion

Experiment 1 showed that the detection of non-adjacent dependencies was hindered when they were embedded in other linguistic materials that were randomly distributed. The experiment also showed that subjects did not learn the adjacent order between the elements, ruling out the possibility that the lack of learning non-adjacent dependency was a result of learning of adjacent dependencies. These findings suggest that the mechanisms when the dependent patterns are presented in isolation (e.g., Gómez, 2002; Peña et al., 2002) may not be engaged in embedded contexts. Or, if available, the parameters of the statistical information they require might be quite different.

However, before accepting this conclusion, it is important to address other explanations for why learning was unsuccessful. One possibility is that something other than the distributional manipulation about the particular materials we used made them more difficult to learn in comparison to Gómez (2002). To test this, we replicated the basic Gómez (2002) design with our materials in Experiment 2.

Experiment 2

In Experiment 1, we found that learners were unable to learn adjacent and non-adjacent dependencies when the dependencies were embedded. In Experiment 2, we ensure that this result was not due to any idiosyncrasies in the materials that instantiated the dependencies, by replicating the non-embedded experimental design and procedure in Gómez (2002). We chose

the version of the experiment which used the same variability at the middle position with Experiment 1 (24 words). We also extended the experiment to test whether learners could learn adjacent order information and non-adjacent dependencies at the same time, as in Experiment 1.

Methods

Participants. A total of twenty-four undergraduate students at University of Southern California in the psychology subject pool participated. Subjects were divided into two counterbalancing groups.

Stimuli. We used the same novel words that we used in Experiment 1.

Design and Procedure. Experiment 2 replicated the design of Experiment 1 in Gómez (2002), with 24 intermediate words. In other words, the critical AXB structures were identical to the ones in Experiment 1, but here they were presented as isolated 3-word sequences, with no buffer words. As in Experiment 1, words occurred every 0.8 seconds within an AXB sequence, and between each sequence there was an additional 0.8 seconds of silence. All the distributional properties of the AXB sequences replicated those in Experiment 1.

There were 10 blocks in the training phase, like in Experiment 1. Each block lasted two and a half minutes, with the entire learning phase being around 25 minutes. The test phase was identical to that in Experiment 1.

Results and Discussion

In Experiment 2, we successfully replicated Gómez (2002), in that participants learned the non-adjacent dependencies. The mean accuracy for the test trials for non-adjacent dependencies was at 71.42% (Figure 2). To test the results statistically, we conducted mixed effects logistic regressions with the binary response (yes/no) as the dependent measure, and counterbalancing group as a between-subjects fixed effect. The results we report here included

by-subject intercepts as a random effect; we also ran models with random item intercepts, but since they did not improve model fit, we do not report them here. The two counterbalancing groups were not statistically different ($\beta=-0.131$, $z=-0.27$, $p=0.787$; n.s.); but the intercept was significant ($\beta=0.976$, $z=2.85$, $p=0.004$), suggesting that the performance was above chance. In addition, participants successfully learned the adjacent orders with the mean accuracy of 75.95% (Figure 1), and the intercept was significantly above chance ($\beta= 0.476$, $z=2.06$, $p=0.039$). There was also a main effect of counterbalancing group ($\beta= 1.424$, $z=3.65$, $p<0.001$), indicating that subjects' accuracy was different across each group. Additional analyses for each group showed that, despite accuracy being different, subjects in both groups were significantly above chance in judging the adjacent sequence orders ($\beta= 0.4575$, $z=2.344$, $p= 0.019$ and $\beta=2.0518$, $z=4.655$, $p<0.001$ for Group 1 & 2, respectively). We therefore found no evidence that attention to the non-adjacent patterns was triggered by difficulty in remembering the adjacent sequences. These findings are consistent with those in Romberg and Saffran (2013), and Vuong, Meyer, and Christiansen (2015), where subjects demonstrated simultaneous learning of adjacent and non-adjacent dependencies.

Considered together with Experiment 1, these results suggest that embedding non-adjacent dependencies in longer utterances hinders successful learning of non-adjacent dependencies, as well as order information between adjacent items. The same patterns that were successfully learned when they were presented in isolation were apparently not-learnable when surrounded by other words (at least with the presentation characteristics in these experiments).

We thus replicated Gómez's (2002) study with 24 intermediate words, alleviating the concern that the null results in Experiment 1 were due to any factors from our procedure or materials. This suggested that embedding non-adjacent dependencies in longer passages in

Experiment 1 interfered with learners' ability to detect them. This set of results is in line with the interpretation that non-adjacent dependencies can only be learned at edge positions.

However, it is also possible that embedding creates a more difficult learning situation in general, for all types of dependencies, and that learning in embedded and non-embedded scenarios is not qualitatively different. This suggests that any dependency in longer sequences would make it more difficult to learn, and it is not a particular case with non-adjacent dependencies.

The results from the test of adjacent dependencies in Experiments 1 and 2 already suggest that adjacent and non-adjacent dependency learning both benefit from being located at the edges of sequences, suggesting similar learning mechanisms. However, we tested adjacent dependency learning in an atypical way, by testing sensitivity to sequential order. As we noted earlier, both element order and class order of the ungrammatical items in the adjacent dependency tests were different from the grammatical items. Furthermore, the transitional probability between the elements was fairly low in the grammatical items (e.g., an AX word would have a forward transitional probability of 1/24). For these reasons, any patterns we observed from adjacent dependency learning in Experiments 1 and 2 cannot be directly compared to the results from non-adjacent dependency learning (where the transitional probability between A and B words in a grammatical test item A_iXB_i is 1). In Experiments 3 and 4, we carry out a variant of Experiments 1 and 2, but specifically probe adjacent dependency learning, and in a more direct way. Across experiments we manipulated whether adjacent dependencies were embedded or not, and assessed the effects of this manipulation on learning. This contrast further informs the discussion of whether NAD learning is qualitatively different from learning adjacent dependencies.

Experiment 3

Results from Experiments 1 and 2 together suggest that learning non-adjacent dependencies is possible when the dependencies are not embedded in longer sequences. Since we are also interested in adjacent dependency learning, we tested learning of adjacent order relations between A, X, and B. But this test did not address whether adjacent dependency learning is subject to similar learning constraints, because learning order information between adjacent items is not the same as learning their dependency. In Experiment 3, we modify the design of Experiment 1 to specifically target adjacent dependencies learning between A and B words. To make it comparable to the previous experiments, we simply altered the position of the X element so that it immediately followed the AB dependency, instead of intervening between A and B positions. In other words, we duplicated the design of Experiment 1 but made the non-adjacent dependency an adjacent one.

Methods

Participants. Twenty-four undergraduate students at University of Southern California recruited from psychology subject pool participated. Subjects were divided equally into two counterbalancing conditions that determined the training language.

Stimulus Creation and Design. We used the same novel words that we used in Experiment 1. Again, word files generated were all shorter than 0.8 s, so silences were added such that concatenating word files results in words occurring every 0.8 seconds within a “sentence”. Between each artificial sentence, there was an additional 0.8 seconds of silence.

The adjacent dependencies in the artificial language were designed by switching AXB pattern from Experiment 1 to the ABX pattern. Similar to Experiment 1, there were three A_iB_i frames and 24 words in the intervening position (i.e., $A_iB_iX_j$, where $i=1-3$ and $j=1-24$). The two

counterbalancing languages were created by taking three pairs of A_B for the three frames in one training language (A_1B_1 , A_2B_2 , A_3B_3), and three different pairs for the other training language (A_1B_2 , A_2B_3 , A_3B_1). This design allowed the use of the same set of test items for the two counterbalancing training languages, where the set included both frames from the two training languages.

All other aspects of the material design were as in Experiment 1.

Procedure. Similar to Experiment 1, there was a learning phase and a testing phase. The procedure for the learning phase was identical to that in Experiment 1. In the testing phase, we used 6 adjacent dependency test items, 3 grammatical and 3 ungrammatical. The grammatical items were taken from the items presented in the learning phase (e.g., $A_1B_1X_1$, $A_2B_2X_1$, $A_3B_3X_1$), while the ungrammatical foils were taken from the non-exposed language (e.g., $A_1B_2X_1$, $A_2B_3X_1$, $A_3B_1X_1$). The only difference between the two test types was whether the adjacent dependency was one that occurred in the familiarized language. Subjects heard testing sequences through headphones, and on each trial answered the question on the screen: “Did you hear this sequence before?” A Y/N keyboard response was collected for this question. After a short pause, the next test trial began. When the study ended, we thanked our participants and debriefed them about the purpose of the study.

Results

The mean accuracy for the test trials, averaging across groups, was at 52.77% (Figure 4). We analyzed performance with the same mixed effects logistic regression model as in the previous experiments. There was no main effect of counterbalancing condition ($\beta=-0.223$, $z=-0.67$, $p=0.505$, n.s.), and the intercept was not statistically different from chance ($\beta=0.223$, $z=0.94$, $p=0.347$, n.s.).

----- Figure 4 Here -----

In sum, participants showed no evidence that they learned adjacent dependencies when they were embedded in the same way we embedded non-adjacent dependencies (Experiment 1).

Experiment 4

In Experiment 3, we found that learners were unable to learn adjacent dependencies when the dependencies were embedded in longer sequences. In Experiment 4, we demonstrate that non-embedded adjacent dependencies are learnable. To achieve this, we replicated the design of Experiment 3, but excluded the buffer words.

Methods

Participants. A total of twenty-four undergraduate students at University of Southern California in the psychology subject pool participated. Subjects were divided into two counterbalancing groups.

Stimuli. We used the same novel words that we used in Experiment 1.

Design and Procedure. Experiment 4 used the ABX design from Experiment 3, where the dependency AB is at the initial position of the sequence with no buffer words. There were 10 blocks in the training phase, like in Experiment 3. Each block lasted two and a half minutes, with the entire learning phase being around 25 minutes. The test phase was identical to that in Experiment 3.

Results and Discussion

The mean accuracy for the test trials, averaging across counterbalancing group, was

82.64% (Figure 4). We analyzed adjacent dependencies with the same mixed effects logistic regressions as the previous experiments. There was no main effect of counterbalancing group ($\beta = -0.461$, $z = -0.58$, $p = 0.562$), but the intercept was above chance, indicating that learning was significantly above chance ($\beta = 2.307$, $z = 3.51$, $p < 0.001$).

Considered together with Experiment 3, these results suggest that embedding adjacent dependencies hinders successful learning of adjacent and non-adjacent dependencies. The same patterns that were successfully learned when they were presented in isolation were apparently not-learnable when surrounded by other words.

These results suggest that learning adjacent dependencies may be subject to the same kinds of learning constraints as learning non-adjacent dependency. Together with the results from Experiment 1 and 2 on non-adjacent dependency, the results from Experiments 3 and 4 suggest the learning mechanisms involved in adjacent and non-adjacent dependency learning may be qualitatively the same, contrary to some prior proposals. We return to this topic in general discussion.

Summary of Experimental Findings

In a series of experiments, we explored the effects of sequential embedding on the learning of adjacent and non-adjacent dependencies. Experiments 1 and 2 set up a contrast between embedded and non-embedded non-adjacent dependency learning, looking at whether non-adjacent dependency learning required being surrounded by pauses. While we also assessed some aspects of adjacent sequence learning in those experiments, Experiments 3 and 4 enabled a direct comparison between learning of adjacent and non-adjacent dependencies under comparable conditions. In the embedding condition where the dependencies were presented amidst other linguistic materials, participants did not show evidence that they learned either the

non-adjacent dependencies (Experiment 1), or the adjacent dependency (Experiment 3). In the contrasting conditions where the dependencies are surrounded by pauses, participants learned both non-adjacent dependency (Experiment 2) and adjacent dependency (Experiment 4). The basic finding was that removing the materials surrounding the dependency changed whether the dependencies were learnable, even though this change did not change the nature or the complexity of the dependencies per se. Moreover, the effect of embedding was similar for adjacent and non-adjacent dependency learning.

General Discussion

In this study, we found that given similar learning situations, learning adjacent and non-adjacent dependencies are subject to the same kinds of learning constraints: Learning is successful when the dependencies occur at the silence-delimited boundaries of relatively short utterances. This finding calls into question theories that claim that NAD learning is carried out by substantially different processes, as those theories relied on there being categorical differences in the kinds of environments that support adjacent and non-adjacent dependency learning (Endress & Bonatti, 2007; Endress et al., 2009; Peña et al., 2002). We now discuss how the present findings can be understood with respect to prior studies.

Our results may appear to be at odds with other studies that report adjacent dependency learning in embedded contexts, such as Saffran et al. (1996) and related studies. However, the findings from Experiment 3 are not incompatible with that literature. Our materials and design differ considerably from those in Saffran et al., and several of these differences may make tracking dependencies in our study more difficult. For one thing, there are simply a greater number of entities to keep track of—72 in our study, versus 12 in Saffran et al.—so the computational complexity is greater. In addition, the rate of speech in Saffran et al. was faster

overall than in our experiments, which has the effect of bringing dependent items in closer temporal proximity, which may facilitate dependency learning (indeed, Wang, Zevin, & Mintz, in prep., show that increasing the speech rate improves NAD learning). Importantly, we are not claiming that adjacent dependencies are unlearnable when they are sequentially embedded in longer passages. Rather, our claim is that when we increase the overall complexity of learning, we see that learning adjacent and non-adjacent dependencies are both facilitated by the same stimulus manipulations. In particular, occurring at the edges of short utterances is not uniquely beneficial to NAD learning. One might even speculate that adding boundaries of silence would further enhance learning in Saffran et al., type materials as well; but since segmentation was the learning criterion, a different test of adjacent dependency learning would have to be used.

One outstanding question is whether the facilitation of dependency learning in Experiments 2 and 4 resulted from the fact that the dependencies occurred at edge positions, or whether it was simply that they occurred in relatively short sequences. Recall from our discussion in the introduction that shorter sequences reduce the possible spans of the dependency—where adjacent dependencies simply involve the most minimal span—which in turn constrains the computational complexity involved in detecting the dependencies (see also Newport, 1990). Edge positions would provide a different kind of benefit, by potentially making the dependent elements more salient. Of course, these factors could interact, such that any benefit deriving from being located at an edge might come into play only for longer sequences. Future research should investigate this question.

Just as we are not claiming that adjacent dependencies are unlearnable when they occur in longer sequences, we are also not claiming that NADs are unlearnable in longer sequences. Some evidence showing that they are learnable come from Frost and Monaghan (2016), who

carried out a version of Peña et al.'s (2002) study, but used more appropriately controlled test items. They found evidence that was consistent with NAD learning, although because of the nature of their test materials, other interpretations were possible. And, as mentioned earlier, a study by Wang et al. (in prep.) showed that increasing the rate of presentation makes embedded NAD learning possible. (It's worth noting that the speech rate in Frost & Monaghan, 2016, was also faster.) Rather, our present findings, and the ones just mentioned, demonstrate that stimulus properties that facilitate NAD learning also seem to facilitate adjacent dependency learning, suggesting that the learning mechanisms are fundamentally the same.

It is also worth noting that there are studies where partially embedded NADs were learned, when the speech rate was similar to the present study (Mintz et al., 2014; Reeder, Newport & Aslin, 2013). It is hard to say how they were learned, since in those artificial languages the dependencies sometimes occurred in relative isolation, and other times they were partially embedded. It is possible that the cases when the dependencies were isolated by pauses were sufficient to support learning. Overall, it could simply be that NAD learning is just harder than learning adjacent dependencies, so that under some conditions, adjacent dependencies are learnable where NAD learning is not, or where it may require additional exposure. Here we made adjacent dependency learning harder—perhaps by slowing down the speech rate—and we showed that edges and shorter utterances benefit adjacent dependency learning, just as those stimulus properties benefit NAD learning. The studies mentioned in the previous paragraph suggest that a faster (and more natural) speech rate might make speech processing easier overall, and thereby facilitate learning NADs. This is a possibility we are currently investigating.

Before we conclude, it is important to note some of the limitations. In order to manipulate whether items are at the edge positions, the experimental condition (embedding) necessarily

brings about some confounding factors with respect to the distributional properties of the sequences used, including the lengths of sentences, the number of tokens used in all sentences, the position where the dependency appears within a sentence, and so on. Some of the confounding factors are unavoidable if one manipulates whether the dependencies are at edge positions (for example, the length of sentences is necessarily longer when dependencies are embedded compared to when they are not embedded). There are other factors, such as the total number of tokens used in the language, that could have been controlled differently, but it would bring about additional confounding factors (for example, making sentences longer without increasing number of token used means repetition of words within sentences). Given these confounding factors, we should be cautious regarding the conclusions we draw from the experimental evidence. For example, we cannot infer the specifics of how non-adjacent dependency is learned from the current experiments (e.g., transitional probabilities, sequence frequency, etc.), and whether the specific type of statistical information is the same when learners compute adjacent and non-adjacent dependencies.

Finally, as a methodological note, it is possible that the testing paradigm adopted here and in many other studies, is not be well suited for testing the kind of statistical learning that may be involved in dependency learning. Statistical learning is often characterized as implicit learning (e.g., Saffran, Newport, Aslin, Tunick, & Barrueco 1997; Turk-Browne, Jungé, & Scholl 2005), and studies on implicit learning suggest that the representations resulting from learning should be probed with implicit measures. Asking yes/no questions about whether particular phrases are in a language requires some explicit representation for the phrases, and this may not be the most relevant way of testing whether implicit learning is successful (Siegelman, Bogaerts, Christiansen, & Frost, 2017). One example of implicit measures includes the use of

SRT task (Vuong, et al., 2015), where reaction time measures are taken as evidence for sequence learning. We are working on incorporating these measures in future research.

Conclusion

The structure of natural languages give rise to many dependencies in the linear sequences of words. Detecting these dependencies is arguably critical for young children in learning the underlying structure of their language. Some prior research has claimed that the kinds of learning mechanisms learners deploy in processing adjacent dependencies are fundamentally different from those deployed in non-adjacent dependency (NAD) learning. Much of the evidence for this claim comes from studies that showed that learners acquire NADs only when the dependencies are isolated by silences in the speech signal, whereas adjacent dependency learning is robust even when the dependencies are embedded in long, continuous sequences of speech. In this study, we challenged this claim by making the stimulus properties and experimental scenario for adjacent and non-adjacent dependency learning highly similar. Under these similar conditions, we found that subjects failed to learn both adjacent and non-adjacent dependencies when they were embedded in longer sequences, and they succeeded at learning when the dependencies occurred in short sequences, isolated by silences. This finding, as well as other recent findings in the literature, challenge the necessity of positing more than one kind of learning mechanism for learning adjacent and non-adjacent dependencies. Rather, we propose that similar mechanisms are involved, but that because NAD learning is more computationally complex, it may be hindered in situations in which adjacent dependency learning nevertheless occurs. Critically, however, the same factors that give a boost to NAD learning, give a boost to adjacent dependency learning as well. We propose that this is because the same kinds of mechanisms are involved.

References

- Baddeley, A. D., Thomson, N., & Buchanan, M. (1975). Word length and the structure of short-term memory. *Journal of Verbal Learning And Verbal Behavior*, 14, 575-589.
- Endress, A. D., & Bonatti, L. L. (2007). Rapid learning of syllable classes from a perceptually continuous speech stream. *Cognition*, 105, 247–299.
<http://doi.org/10.1016/j.cognition.2006.09.010>
- Endress, A.D., Nespors, M. & Mehler, J. (2009). Perceptual and memory constraints on language acquisition. *Trends in Cognitive Sciences*, 13, 348-353.
- Gallistel, C. R. (2009). The importance of proving the null. *Psychological review*, 116, 439.
- Gómez, R. L. (2002). Variability and detection of invariant structure. *Psychological Science*, 13, 431-436.
- Hauser, M. D., Newport, E. L., & Aslin, R. N. (2001). Segmentation of the *speech* stream in a non-human primate: statistical learning in cotton-top tamarins. *Cognition*, 78, B53-B64.
- 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.
- Peña, M., Bonatti, L. L., Nespors, M., & Mehler, J. (2002). Signal-driven computations in speech processing. *Science*, 298, 604-607.
- 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.
- Saffran, J. R., Newport, E. L., Aslin, R. N., Tunick, R. A., & Barrueco, S. (1997). Incidental Language Learning: Listening (And Learning) out of the Corner of Your Ear. *Psychological Science*, 8, 101-105.
- Santelmann, L. M., & Jusczyk, P. W. (1998). Sensitivity to discontinuous dependencies in language learners: Evidence for limitations in processing space. *Cognition*, 69, 105-134.
- Siegelman, N., Bogaerts, L., Christiansen, M. H., & Frost, R. (2017). Towards a theory of individual differences in statistical learning. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 372, 20160059.
<http://doi.org/10.1098/rstb.2016.0059>
- Swingley, D. (2005). Statistical clustering and the contents of the infant vocabulary. *Cognitive Psychology*, 50, 86-132.
- Toro, J. M., & Trobalón, J. B. (2005). Statistical computations over a speech stream in a rodent. *Attention, Perception, & Psychophysics*, 67, 867-875.
- Turk-Browne, N. B., Jungé, J. A., & Scholl, B. J. (2005). The automaticity of visual statistical learning. *Journal of Experimental Psychology: General*, 134, 552-564.
- Van den Bos, E., & Christiansen, M.H. (2009). Sensitivity to non-adjacent dependencies embedded in sequences of symbols. *Proceedings of the 31st Annual Meeting of the Cognitive Science Society*.
- Vuong, L. C., Meyer, A. S., & Christiansen, M. H. (2015). Concurrent Statistical Learning of Adjacent and Nonadjacent Dependencies. *Language Learning*, 66, 8-30.
<http://doi.org/10.1111/lang.12137>

Wang, F. H., Zevin, J. D., & Mintz, T. H. (in prep). Learning Non-Adjacent Dependencies in Continuous Presentation of an Artificial Language.

Appendix: Example Sequences from Experiments

The examples below are to show some of the possible sequences in the experiments using one set of non-adjacent dependency (*choon _ jub*); words making up the dependency are italicized and underlined for demonstration purposes only: none of the auditory stimuli of the words in the experiments are marked in any way. Recall that each word is padded at the end with silence such that the word plus silence duration is 0.8 s. This means that there are brief, variable length silences between words within a sentence, in addition to the longer between-sentence pauses that are shown.

Experiment 1: blit *choon* balip *jub* ghire paziv (0.8 s pause) channer lowfa *choon* coomo *jub* ghen gip refton (0.8 s pause) swech klide lific *choon* deecha *jub* dess (0.8 s pause) ...

Experiment 2: *choon* balip *jub* (0.8 s pause) *choon* coomo *jub* (0.8 s pause) *choon* deecha *jub* (0.8 s pause) ...

Experiment 3: blit *choon jub* balip ghire paziv (0.8 s pause) channer lowfa *choon jub* coomo ghen gip refton (0.8 s pause) swech klide lific *choon jub* deecha dess (0.8 s pause) ...

Experiment 4: *choon jub* balip (0.8 s pause) *choon jub* coomo (0.8 s pause) *choon jub* deecha (0.8 s pause) ...

Table 1. List of Words used in Experiment 1-4

A words	choon, nud, tiv
B words	jub, dap, fex
X words	antow, balip, benez, coomo, deecha, fengle, gensim, glaik, gople, hiftam, kicey, lartsue, loga, malsig, mirshow, plizit, pooda, puser, roosa, skiger, skyteb, suleb, vamey, wadim
I and F words	blit, channer, dess, ghen, ghire, gip, grimpot, klide, lific, lowfa, paziv, plosit, refton, sowch, swech, voy

Table 2. Illustration of grammar used to sequence words in Experiment 1.

Grammar rules	Explanation & Examples
$S \rightarrow I A X B F$	Five parts are required to make up a sentence (S)
$I \rightarrow i_1$	Sentence initial position (I) can be 1 word, 2 words or 3 words
$I \rightarrow i_1 i_2$	Examples: blit; channer lowfa; swech klide lific
$I \rightarrow i_1 i_2 i_3$	
$A \rightarrow a_n$	A position contains only one word (A_1 , A_2 or A_3)
	Examples: choon; nud; tiv
$X \rightarrow x_m$	X position contains only one word (X_1 through X_{24})
	Examples: balip; benez; coomo; deecha
$B \rightarrow b_n$	B position contains only one word (B_1 , B_2 or B_3)
	Examples: jub; dap; fex
$F \rightarrow f_1$	Sentence final position (F) can be 1 word, 2 words or 3 words
$F \rightarrow f_1 f_2$	Examples: dess; ghire paziv; ghen gip refton
$F \rightarrow f_1 f_2 f_3$

Figure 1. Performance on adjacent order knowledge test in Experiment 1 & 2. Error bars depict 95% confidence intervals calculated in logit units on the corresponding β 's for the logistic regressions that compared outcomes to chance; logits were transformed to percentages for this graph.

Figure 2. Tests of non-adjacent dependencies in Experiment 1 & 2. Error bars depict 95% confidence intervals calculated in logit units on the corresponding β 's for the logistic regressions that compared outcomes to chance; logits were transformed to percentages for this graph.

Figure 4. Tests of adjacent dependencies in Experiment 3 & 4. Error bars depict 95% confidence intervals calculated in logit units on the corresponding β 's for the logistic regressions that compared outcomes to chance; logits were transformed to percentages for this graph.

Figure 3. Bayes factors at both group level and individual subject level. In graphs A & B for group performance, each point represents the odds (Bayes Factor) in favor of the assumed null prior (y-axis) for a particular range of alternative probabilities (x-axis). Graph A shows the results for adjacent order tests, and Graph B shows the results for non-adjacent dependencies tests. For these two graphs (using the analysis strategy from Gallistel, 2009): the value on the x-axis represented the upper bound on a uniform probability distribution for the alternative prior starting at 0.5; hence, the vagueness of the alternative prior increased from left to right. The y-axis is plotted on a log scale, where values above 2^0 (1) suggest evidence for the null, and below 2^0 suggested evidence against the null. Individual results are in Graph C & D, where every two bars represented the data from one participant. Graph C shows the results for adjacent order tests,

and Graph D shows the results for non-adjacent dependencies tests. For these two graphs, the left bar represents the results when the alternative distribution ranges from 0.5 to 0.65, and the right bar from 0.5 to 1. Data for individual subjects are plotted in the same ordinal position on these latter two graphs (e.g., the first two bars from the top graph and the bottom graph came from the same participant).