Replication of Saffran, Newport, & Aslin (1996) Word segmentation: The role of distributional cues, Exp. 1

Joshua K. Hartshorne

Department of Psychology, Boston College

Contact joshua.hartshorne@bc.edu

I replicated Exp. 1 of Saffran, Newport, & Aslin (1996) Word segmentation: The role of distributional cues, *Journal of Memory and Language*, 35, 606-621, after a prior, largely unsuccessful replication attempt (Garcia et al., 2017, Replication of Saffran, Newport, & Aslin (1996) Word segmentation: The role of distributional cues, Exp. 1. PsyArXiv doi:10.17605/OSF.IO/QSYD2). The present replication corrected a randomization error in the prior replication and introduced attention checks to ensure that the participants were indeed attending to the stimuli. Despite this, and despite a much larger sample than the original (100 subjects vs. 24), evidence of successful segmentation was once again weak and mixed, and none of the item or condition effects replicated. I consider whether this is more likely to be a failure of replication or a failure of generalization (e.g., to a different population).

Keywords: statistical learning, unsupervised learning, language acquisition, replication

In this paper, I describe a replication of Experiment 1 from a classic statistical word-segmentation paper by Saffran, Newport, and Aslin 1996 (henceforth: SNA96). This experiment was part of a larger project to systematically replicate as many experiments involving statistical word learning in adults as possible.

Method

Subjects

100 individuals participated in exchange for monetary compensation, 98 of whom were native English speakers. These included 50 subjects in the part-word condition (Ages 20-58, $M=32$) and 50 subjects in the non-word condition (Ages 20-60, $M=33$).

Known differences with original.

- Our study had an order of magnitude more subjects (24 vs. 100) and thus considerably more power and precision.
- Subjects in our study were significantly older on average than the original subjects, who were all undergraduates.
- Subjects in our study were compensated monetarily rather than with course credit.

Materials

The language consisted of four consonants (/p/, /t/, /b/, /d/) and three vowels (/a/, /i/, /u/) which, when combined,

rendered an inventory of 12 CV syllables. 11 of these syllables (excluding /di/) were combined to create six trisyllabic words: babupu, bupada, dutaba, patubi, pidabu, and tutibu. Some members of the syllable inventory occur in more words than others; for example, /bu/ occurs in four words, while /ta/ occurs in only one. This ensures varying transitional probabilities within the words themselves

Each syllable was generated using MBROLA with voice *us1*. Each syllable was produced in context with full coarticulation between syllables. For each syllable, the consonant lasted for 100 ms and the vowel for 177 ms, for a rate of 216 syllables/minute. 252 tokens of each of the 6 words were concatenated in a random order, with the stipulation that the same word never occurred twice in a row. The resulting sound file was converted from WAV to MP3 using SoX.

Average transition probability within a word ranged from 42.0% to 100%. Transition probabilities between phonemes spanning words ranged from 5.0% to 59.0%.

For the test phase, six nonword foils and six part-word foils were created. The nonwords (*pibuda*, *badapu*, *tapubi*, *bubita*, *dubiti*, *bubibi*) consisted of syllables from the language's syllable inventory which never followed each other in the speech stream, even across word boundaries. Thus, the transitional probabilities between each of the syllables in the nonwords were zero. Each of the six part-words consisted of a syllable pair from a word from the language, plus an additional syllable. For example, the final syllable of the word *pidabu* was altered to create the part-word *pidata*. Similarly, the first syllable of the word *dutaba* was altered to create the part-word *bitaba*. Three part-words contained the first two

syllables of words (*pidata*, *bupabi*, *babuda*) and three contained the final two syllables of words (*bitaba*, *datubi*, *titibu*).

The stimuli for the test phase were generated from the same syllables used to create the training, using the procedure described above. 72 test trials were created by combining each real word with each foil, with an approximately 500 ms gap in between with the stipulation that the real word appeared first in half of the trials in each condition (nonword or partword).

In addition, we generated 6 catch trials, which pitted each word against a (different) foil that contained multiple phonemes that did not appear in the training.

Known differences with original.

- The original stimuli were produced with MacinTalk. We used MBROLA.
- We used all three of the partwords that are explicitly mentioned in the original paper (*pidata*, *bitaba*, *bupabi*). Otherwise, our foils likely differ from theirs.
- The original paper states that training included 300 tokens of each of the 6 words at 216 syllables/minute for a total of 21 minutes. These numbers are mutually incompatible: Either the speed was faster (257 syllables/minute), the length of the audio was longer (25 min), or there were fewer tokens (252). The latter number matches the assertion elsewhere in the same paragraph that there were a total of 4,536 syllables. Thus, we presented 252 tokens of each of the 6 words in order to have 21 minutes at 216 syllables/minute.
- The original paper reports that the transition probabilities between phonemes spanning words ranges from 10% to 20%. This is inconsistent with the rest of the reported method. If they constructed training as described, the transition probability between the words *dutaba* and *bupada* would have to be at least 50%. This follows from that fact that in at least 50% of the cases where /ba/ appears, it is within the word *babupu*. Thus, if *bupada* ever followed *dutaba* in training, the transitional probability of /babu/ would be greater than 50%. The probability of that not having happened is negligible ($p < 4 * 10^{-25}$). For similar reasons, the transitional probability between the words *pidabu* and *patubi* would be around 35%. The transition probabilities between phonemes spanning words in our experiment ranged from 5.0% to 59.0%.
- The average transition probability within-word were 42.0%, 42.5%, 50.0%, 75.0%, 75.0%, and 100%. The original paper reports 37%, 41%, 50%, 75%, 75%, and 100%. The difference could be due to differences in the random orders of the stimuli or due to a miscalculation on the part of the original authors.

Possible differences with original.

- In our version, all six of the real words were used to create partwords. The original paper does not specify whether all words were used to create partwords.
- In our version, we balanced the number of test trials where the correct answer came first or second. No information regarding this issue is provided in the original paper.

Apparatus

The experiment was programmed using jsPsych (de Leeuw, 2014) and run on Amazon Mechanical Turk using Gureckis et al. (2016). As the experiment was conducted online, we were unable to control the listening environment of the participants on Mechanical Turk, but instructed them before the experiment to sit in a quiet and non-distracting environment with the sound turned up to a sufficient level.

Known differences with original. The original experiment was conducted in an IAC sound-attenuated booth. Stimuli were presented using an Aiwa tape deck and a Proton speaker. The test phase was administered on a Macintosh IIsx computer, using the PsychLab stimulus presentation package.

Procedure

Subjects were instructed to listen to a "nonsense" language. They were told that the language contained words, but no meanings or grammar. They were informed that their task was to figure out where the words began and ended. Subjects were given no information about the length or structure of the words or how many words the language contained. They were informed that the listening phase of the experiment consisted of three short blocks, followed by a test of their knowledge of the words in the language. Subjects were allowed to take a break after each of the first two 7-min listening blocks.

After a total of 21 min of listening, subjects received a two-alternative forced-choice test. For each test item, subjects heard two trisyllabic strings, separated by 500 ms of silence. One of these strings was a word from the nonsense language, while the other was not. Subjects were asked to indicate which of the two strings sounded more like a word from the language by pressing either the "1" or "2" key on the keyboard. The test items were constructed by pairing the six words of the language with one of two types of trisyllabic foils. For half of the subjects ($n=50$), the six nonword foils were illegal strings of syllables from the language. For the other half of the subjects ($n=50$), the six foils were partwords, consisting of a syllable pair from a word plus an additional syllable. After the subject's response, the next item was presented. The order of presentation of items was randomized for each subject.

After the primary test trials, all subjects answered the 6 catch trials. The order was randomized for each subject. The catch trials followed the primary test trials immediately, without any overt distinction.

Known differences with original.

- The original study enforced a five-minute break between training blocks. Given the nature of online testing, this was difficult to enforce, so we allowed subjects to self-pace their breaks.
- The original provided practice trials were given to each subject prior to the test, in order to clarify the structure of the test and to give subjects practice with the key presses. The practice trials consisted of strings of syllables which were not part of the language's syllable inventory. We did not include practice trials, as it seemed unlikely that participants in an online study were unfamiliar with computers.
- The original study did not include catch trials. Note that because the catch trials follow the primary test trials, they cannot affect the results of those trials (except in that we can use the catch trials to distinguish poor learners from noncompliant subjects).

Results

All data and code are available through the Open Science Framework at osf.io/bj4cs.

All Subjects

We first considered results from all subjects, regardless of performance on catch trials. This matches the analyses of Garcia et al. Garcia et al..

Original Analyses. The overall results are presented in Fig. 1. In the nonword test condition, the mean score was 21.8 of a possible 36 (61%, SE=2%), where chance performance equals 18. A single-sample t-test (two-tailed) showed overall performance significantly different from chance: $t(49)=5.80$, $p=4.7 \times 10^{-7}$. Four of the six words (*bupada*, *patubi*, *pidabu*, *tutibu*) in this condition were learned at a level significantly better than would be expected by chance ($p < 0.05$ for each word).¹ Although the foils in the partword condition were designed to be more confusable with the words, the mean score for participants was still 20.3 of a possible 36 (57%, SE=1%), a rate that significantly exceeded chance ($t(49)=4.52$, $p=3.9 \times 10^{-5}$). However, success was not significantly different from that in the nonword condition ($t(98)=1.77$, $p=.08$). Three of the six words (*dutaba*, *patubi*, *pidabu*) were learned significantly better than would be expected by chance ($ps < .01$). **Thus, we replicated the finding of statistical word segmentation, albeit at a much lower rate.**

all subjects

attentive subjects

Figure 1. Histograms of subject accuracy in the nonword condition (left) and partword condition (right), considering all subjects (top) or only those with perfect accuracy on catch trials (bottom).

Next, we asked whether the strengths of the transitional probabilities played a role in which words were learned better than others. Recall that the transitional probabilities between the pairs of syllables within words ranged between .43 and 1. We split the six words into two sets, one set containing the three words with the highest average transitional probabilities (1.0, .77, and .75), and the other set containing the three words with the lowest average transitional probabilities (.50, .43, and .43). A matched-pairs t-test comparing subjects' mean performance on the three words with the highest transitional probabilities ($M=58\%$, $SE=1\%$) to the three words with the lowest transitional probabilities ($M=59\%$, $SE=1\%$) indicated no significant difference in performance on the words containing higher transitional probabilities: $t(99)=0.26$, $p=0.79$. **Thus, we do not replicate the prior finding that words with higher transition probabilities were more easily recognized.**

SNA96 report that accuracy was particularly high on partword trials with the foil *bupabi*. This is relevant because dur-

¹SNA96 do not state what statistical test was applied to the data for individual words. Because they used t-tests elsewhere, we assume they did so here as well.

ing training, the syllable pair /bupa/ was more common than any other. If subjects were simply following bigram frequencies, they should have relatively poor performance on this item. Unlike SNA96, that is what we found. Performance on such trials was low ($M=52\%$, $SE=3\%$) and indistinguishable from chance ($t(49)=0.54$, $p=.59$). **Thus, we were unable to replicate this finding.**

These results may be probed further to determine which aspects of the words were learned. Recall that the part-word test used two types of part-words. Half of the part-words contained the first two syllables of words plus a final syllable, and the other half contained the final two syllables of words plus an initial syllable. SNA96 found that while the first type produced false alarm rates significantly lower than chance ($M=29\%$; $t(11)=-61.3$, $p<.01$), the latter do not ($M=48\%$; $t(11)<1$), a difference that reached significance ($t(11)=3.25$, $p<.01$). Our prior replication found the exact opposite (.46 vs. .38; $t(49)=3.56$, $p=.001$). In contrast, the present study reveals below-chance performance for both change-final foils ($M=43\%$, $t(49)=3.60$, $p=.0007$) and change-first foils ($M=44\%$, $t(49)=3.24$, $p=.002$), with no significant difference between the two ($t(49)=0.67$, $p=.50$). **Thus, we were unable to replicate this finding.** All three studies used different foils, suggesting that the effects for this analysis may be driven by the specific foils themselves.

Revised Analyses

There are two important limitations to the original analyses conducted by SNA96. First, their use of t-tests on accuracy data violates core assumptions of t-tests: that the data are continuous and error is normally distributed. Violation of these assumptions can lead both to false positives and false negatives (Jaeger (2008)). Binomial models are more appropriate (Jaeger (2008)).²

Additionally, they treat items as fixed effects. Thus, any conclusions are specific to the specific stimuli used. No assessment of generalization can be made. Statisticians have suggested two options. The first is to conduct analyses separately over subject averages (as SNA96 did) and over item averages (Clark (1973)). The second option is to use mixed effects models (Baayen et al. (2008)).

Given these considerations, we re-analyzed the data used in binomial mixed effects models, implemented in R using the lme4 package, ver. 1.1-12 (Bates et al., 2017). Our first model had a single fixed effect: foil type (nonword/part-word). We considered two possible random effects structures. The first estimated random intercepts for each word/foil pair and each subject. The second estimated separate random intercepts for each word and foil separately, random intercepts for subject, and a random subjects with respect to condition. The former provided a better fit as assessed by conditional AIC (8374 v. 8377) (We calculated cAIC using Stein formulas, as implemented in cAIC4 v. 2 (Grevén &

Kneib, 2010; Saefken & Ruegamer, 2015)). Thus, we use the former random effects structure throughout, including random slopes where appropriate.

Thus, we assessed the effect of foil type with binomial mixed effects model with a fixed effect of foil type and random effects as described above. The main effect of foil type was not significant (Wald's $z=-1.45$, $p=.15$).³

We further assessed evidence for learning by running the above model with no fixed effects. The intercept was significant whether collapsing across conditions (Wald's $z=5.76$, $p=8.3 \times 10^{-9}$), or when considering the nonword and part-word condition separately (respectively: $z=4.73$, $p=2.2 \times 10^{-6}$; $z=3.41$, $p=.0007$), indicating that accuracy was significantly different from chance.

Next, we asked whether the three words with higher transition probability were recognized at a different rate from the three with lower transition probability. For this model, we included a fixed effect of transition probability and random effects as described above. The main effect of transition probability was not significant (Wald's $z=0.18$, $p=.86$).

We then determined whether performance on *bupabi* was greater than chance. Again, we found that it was not (Wald's $z=0.4$, $p=.66$).

Finally, we compared false alarms to the partwords that matched real words on the first two syllables against accuracy for those that matched real words on the last two syllables with a model with a fixed effect for partword type and random effects as described above. The main effect of partword type was not significant (Wald's $z=0.52$, $p=.60$).

Subjects who passed catch trials

In order to confirm that the results above were not due to subjects who were not paying attention, we re-ran the above analyses, excluding the 47 subjects who missed *any* catch trials. Note that the catch trials were relatively subtle, so this is a strong filter. This filter resulted in 34 subjects in the nonword condition and 19 in the partword condition. One of the remaining subject was not a native English speaker.

Original Analyses. The overall results are presented in Fig. 1. In the nonword test condition, the mean score was 23.4 of a possible 36 (65%, $SE=2\%$), where chance performance equals 18. A single-sample t-test (two-tailed) showed overall performance significantly different from chance: $t(33)=7.97$, $p=3.4 \times 10^{-9}$. Five of the six words (*bupada*, *dutaba*, *patubi*, *pidabu*, *tutibu*) in this condition were learned at a level significantly better than would be

²An intermediate option is to transform accuracy measures to log-odds, though this method is less accurate than a binomial model.

³Wald's z is somewhat anti-conservative when sample size is small (Luke (2016)). Thus, these analyses may if anything overstate the evidence against the null.

expected by chance ($p < 0.05$ for each word).⁴ Although the foils in the part-word condition were designed to be more confusable with the words, the mean score for participants was 22.2 of a possible 36 (62%, $SE = 2\%$), a rate that significantly exceeded chance ($t(18) = 6.03, p = 1.1 \times 10^{-5}$). However, success was not significantly different from that in the non-word condition ($t(51) = 1.22, p = .23$). Four of the six words (babupu, dutaba, patubi, pidabu) were learned significantly better than would be expected by chance ($ps < .05$). **Thus, we again replicate the prior finding of learning, albeit at a much lower rate.**

Next, we asked whether the strengths of the transitional probabilities played a role in which words were learned better than others. We split the six words into 'high transition probability' and 'low transition probability' sets, as described in the previous section. A matched-pairs t-test comparing subjects' mean performance on the three words with the highest transitional probabilities ($M = 64\%$, $SE = 2\%$) to the three words with the lowest transitional probabilities ($M = 64\%$, $SE = 2\%$) indicated no significant difference in performance on the words containing higher transitional probabilities: $t(52) = 0.18, p = 0.86$. **Thus, we again do not replicate the prior finding that words with higher transition probabilities were more easily recognized.**

As noted above, SNA96 report that accuracy was particularly high on part-word trials with the foil *bupabi*. The theoretical importance is described in the previous section. In contrast, we found that performance on such trials was low ($M = 57\%$, $SE = 5\%$) and indistinguishable from chance ($t(18) = 1.29, p = .21$). **Thus, we again were unable to replicate this finding.**

As described in the previous section, we compared the part-words with a changed first syllable to those with a changed final syllable. While SNA96 found a difference, we again did not. In contrast, the present study reveals below-chance false alarms for both change-final foils ($M = 36\%$, $SE = 3\%$, $t(18) = 4.90, p = .0001$) and change-first foils ($M = 41\%$, $SE = 3\%$, $t(18) = 2.88, p = .01$), with no significant difference between the two ($t(18) = 1.21, p = .24$). All three studies used different foils, suggesting that the effects for this analysis may be driven by the specific foils themselves.

Revised Analyses

We again re-analyzed the results using the more appropriate binomial mixed effects models. Our first model had a single fixed effect: foil type (nonword/partword). We considered two possible random effects structures. The first estimated random intercepts for each word/foil pair. The second estimated separate random intercepts for each word and foil separately, plus a random slope for each word. The latter provided a better fit as assessed by conditional AIC (4776 v. 4769) (We calculated cAIC using Stein formulas, as im-

plemented in cAIC4 v. 2 Greven & Kneib, 2010; Saefken & Ruegamer, 2015). Thus, we use the latter random effects structure throughout, including random slopes where appropriate.

Thus, we assessed the effect of foil type with binomial mixed effects model with a fixed effect of foil type and random effects as described above. The main effect of foil type was not significant (Wald's $z = 0.50, p = .61$).⁵

We further assessed evidence for learning by running the above model with no fixed effects. The intercept was significant when collapsing across conditions (Wald's $z = 3.87, p = .0001$), or when considering the nonword and partword condition separately (respectively: $z = 3.39, p = .0007$; $z = 2.13, p = .03$), indicating that accuracy was significantly different from chance.

Next, we asked whether the three words with higher transition probability were recognized at a different rate from the three with lower transition probability. For this model, we included a fixed effect of transition probability and random effects as described above. Unfortunately, the model did not converge, so no conclusions can be made.⁶

We then determined whether performance on *bupabi* was greater than chance. Again, we found that it was not (Wald's $z = 1.0, p = .34$).

Finally, we compared false alarms to the partwords that matched real words on the first two syllables against accuracy for those that matched real words on the last two syllables with a model with a fixed effect for partword type and random effects as described above. The main effect of partword type was not significant (Wald's $z = 0.74, p = .46$).

Discussion

We were able to replicate SNA96's finding of above-chance recognition of words in a statistical learning paradigm. Unlike in our previous replication, we found statistical evidence that this learning would generalize to other stimuli – a question that SNA96 did not address.

Otherwise, we were unable to confirm any of SNA96's other observations. Moreover, learning was considerably weaker than what was reported in SNA96. This remained the case even when we excluded subjects who did not pass catch trials, something that SNA96 did not do. Note that even with

⁴SNA96 do not state what statistical test was applied to the data for individual words. Because they used t-tests elsewhere, we assume they did so here as well.

⁵Wald's z is somewhat anti-conservative when sample size is small Luke (2016). Thus, these analyses may if anything overstate the evidence against the null.

⁶We also ran the analyses filtering out subjects who missed more than one item. In this case, the model did converge. The effect of transition probability was not significant (Wald's $z = 0.0, p = .99$). Results for the other analyses reported in this section were qualitatively similar to what is reported in this section.

that restricted subset of subjects, we had considerably more statistical power than SNA96.

Note that we used a strict criterion for replication: significant results in the same direction. This is both for simplicity and because SNA96 do not provide the information needed for more sophisticated measures of replication (e.g., standard errors).

As with any replication, our replication was inexact. This is particularly true in our case, because of uncertainty as to what the original method was (see above). There are several known differences between our experiment and the original that could have impacted the results. We used different foils, and it may be that ours were *a priori* more difficult to distinguish from the words. Indeed, comparison across our two replications suggests that differences in foils introduce substantial variability. This is perhaps not surprising given the very small number of stimuli. This would be especially likely if the foils in the original used the syllable /di/, which was part of the language but did not appear in the training (see above). Second, subject population was on average older and likely of lower socioeconomic status than SNA96's. Moreover, the fact that our subjects were tested over the Internet could result in them being less attentive. However, when we restricted to subjects who passed the catch trials, this only slightly improved accuracy and did not change the pattern of replication.

References

- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59(4), 390–412.
- Bates, D., Maeschler, M., Bolker, B., Walker, S., Christensen, R. H. B., Singmann, H., ... Green, P. (2017). *Linear mixed-effects models using 'eigen' and 's4'*. cran.r-project.org/web/packages/lme4/index.html.
- Clark, H. H. (1973). The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior*, 12(4), 335–359.
- de Leeuw, J. R. (2014). jspsych: A javascript library for creating behavioral experiments in a web browser. *Behavior Research Methods*, 47(1), 1–12.
- Garcia, C. R., Iozzo, G. L., Lamirato, K. E., Ledoux, J. R., Mu, J., Murdock, K. N., ... Hartshorne, J. K. (2017). Replication of saffran, newport, & aslin (1996) word segmentation: The role of distributional cues, exp. 1. doi: 10.17605/OSF.IO/QSYD2
- Greven, S., & Kneib, T. (2010). On the behaviour of marginal and conditional aic in linear mixed models. *Biometrika*, 773–789.
- Gureckis, T. M., Martin, J., McDonnell, J., Rich, A. S., Markant, D., Coenen, A., ... Chan, P. (2016). psiTurk: An open-source framework for conducting replicable behavioral experiments online. *Behavior Research Methods*, 48(3), 829–842.
- Jaeger, T. F. (2008). Categorical data analysis: Away from ANOVAs (transformation or not) and towards logit mixed models. *Journal of Memory and Language*, 59(4), 434–446. doi: 10.1016/j.jml.2007.11.007
- Luke, S. G. (2016). Evaluating significance in linear mixed-effects models in r. *Behavior Research Methods*, 1–9.
- Saefken, B., & Ruegamer, D. (2015). *Conditional akaike information criterion for lme4*. cran.r-project.org/web/packages/cAIC4/.
- Saffran, J. R., Newport, E. L., & Aslin, R. N. (1996). Word segmentation: The role of distributional cues. *Journal of memory and language*, 35(4), 606–621.