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

Caitlin R. Garcia, Gina L. Iozzo, Katie E. Lamirato, James R. Ledoux, Jesse Mu, Kara N Murdock, Alyssa A. Savery, James E. Spizzirro, Kelsey A. Trimm, Kendall D. van Horne, Juliani Vidal, Sven L. Dietz, Jon Ravid, Lauren Skorb, and Joshua K. Hartshorne
Department of Psychology, Boston College

We replicated Exp. 1 of Saffran, Newport, & Aslin (1996) Word segmentation: The role of distributional cues, *Journal of Memory and Language*, 35, 606-621. This experiment was part of a larger project to systematically replicate as many experiments involving statistical word learning in adults as possible.

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

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

Method

Subjects

101 native English speakers participated in exchange for monetary compensation. These included 50 subjects in the part-word condition (Ages 21-56, M=33) and 51 subjects in the non-word condition (Ages 20-59, M=36). The additional subject in the non-word condition was due to software error.

Known differences with original.

- Our study had an order of magnitude more subjects (24 vs. 101) 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 IBM Watson's Text to Speech with the American English voice *Allison* (IBM, 2017). Each syllable was produced in isolation without co-articulation using the following SSML code. (In test runs Watson TTS tended to slow down at the beginning and end of an utterance, so we included some irrelevant syllables at the beginning and end.)

```
<speech><voice-transformation type="
  Custom" pitch="low"
  pitch_range="x-narrow" rate="x-fast">
  la la
  la. bah. bee. boo. dah. dee. doo. pah.
  pee. poo. tah. tee. too. la la la
.</voice-transformation></speech>
```

The resulting WAV-format sound file was then parsed using Audacity (Team, 2017) into individual syllables. The length of each syllable sound file was 0.227 sec, for a rate of 216 syllables/minute. These syllables were then concatenated to make the six words of the language using Sound eXchange (SoX) (Bagwell & Contributors, 2015). Again using SoX, 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.

Average transition probability within a word ranged from 42.5% to 100%. Transition probabilities between phonemes spanning words ranged from 5% to 60%.¹

This work was supported by a TAME grant from Boston College.

¹These are expected transition probabilities. Because training stimuli were generated using a randomization procedure, the asac-

For the test phase, six nonword foils and six part-word foils were created. The nonwords (*babita*, *dabibi*, *dudata*, *pu data*, *tadupu*, *tutapi*) 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 (*babubi*, *pidata*, *patubu*) and three contained the final two syllables of words (*bitaba*, *budabu*, *tadabu*).

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).²

Known differences with original.

- The original stimuli were produced with MacinTalk with coarticulation between syllables. Our stimuli had no co-articulation between syllables.
- We used both of the foils that are explicitly mentioned in the original paper's Method section (the partwords *pidata*, *bitaba*). Otherwise, our foils likely differ from theirs.³
- 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% to 60%.

- The average transition probability within-word ranged from 42.5% to 100%, whereas the original paper reports a range of 31% to 100%. This likely is due to differences in the randomization (see also earlier footnote).

Possible differences with original.

- In our version, only five of the real words were used to create partwords (*pidabu* was the source for both *budabu* and *tadabu*). 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 psiTurk (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

tual transition probabilities likely differed slightly. Unfortunately, the actual transition probabilities were not recorded, though they could be reconstructed from the audio.

²The exact length varied slightly because the 227 ms sound files for each syllable contained a variable amount of silence at the beginning and end. See above.

³An additional partword, *bupabi*, is described in the Results section, but we did not notice this prior to running the study.

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 ($n=51$), 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 part-words, 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 the same for each subject.⁴

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.
- The order of test trials was the same for each subject.

Results

Original Analyses

All data and code are available through the Open Science Framework at osf.io/z7djs (DOI 10.17605/OSF.IO/Z7DJS).

The overall results are presented in Fig. 1. In the nonword test condition, the mean score was 20.7 of a possible 36 (57%, $SE=2\%$), where chance performance equals 18. A single-sample t-test (two-tailed) showed overall performance significantly different from chance: $t(50)=4.92$, $p=.00001$. Three of the six words (*dutaba*, *patubi*, *tutibu*) in this condition were learned at a level significantly better than would be expected by chance ($p<0.05$ for each word).⁵ Although the foils in the part-word condition were designed to be more confusable with the words, the mean score for participants was 20.9 of a possible 36 (58%, $SE=1\%$), a rate that significantly exceeded chance ($t(49)=6.45$, $p=.00000005$). However, success was not significantly different from that in the

Figure 1. Mean scores for each subject in the two conditions of Experiment 1, of a possible 36 correct.

nonword condition ($t(99)=0.33$, $p=.74$). Only two of the six words (*dutaba*, *patubi*) 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.**

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 ($M=57\%$, $SE=1\%$) to the three words with the lowest transitional probabilities ($M=58\%$, $SE=1\%$) indicated no significant difference in performance on the words containing higher transitional probabilities: $t(100)=0.41$, $p=0.68$. **Thus, we do not replicate the prior finding that words with higher transition probabilities were more easily recognized.**

SNA96 include in their paper a verification that subject performance in distinguishing real words and avoiding foils was not mediated by different levels of syllable co-occurrence in the listening data. Specifically, since the syllable pair /bupa/ occurred more often than other syllable pairs in the training, they verified that performance in avoiding the *bupabi* part-word foil was not unusual compared to the other

⁴The order was intended to be randomized separately for each subject, but was not due to a programmer error.

⁵SNA96 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.

part-word foils. In this replication, we do not have the word *bupabi*, so we do not conduct this analysis.

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 ($t(11)=3.25$, $p<.01$). We found the exact opposite ($M=46\%$, $SE=2\%$ vs. $M=38\%$, $SE=2\%$; $t(49)=3.56$, $p=.001$). **Thus, we were unable to replicate this finding.**

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)). Logit models are more appropriate (Jaeger (2008)).⁶

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. 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 (8441 v. 8448) (We calculated cAIC using Stein formulas, as implemented in cAIC4 v. 2 (Greven & Kneib, 2010; Saeften & 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 $z=1.38$, $p=.17$).⁷

We further assessed evidence for learning by running the above model with no fixed effects. The intercept was non-significant whether collapsing across conditions (Wald's $z=1.69$, $p=.09$), or when considering the nonword and part-word condition separately (respectively: $z=1.35$, $p=.18$;

$z=1.50$, $p=.14$), indicating that accuracy was not 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 $z=1.15$, $p=.25$).

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 $z=0.63$, $p=.53$).

Discussion

We were able to replicate SNA96's finding of above-chance recognition of words in a statistical learning paradigm. However, we are not able to conclude that this learning would generalize to other stimuli – a question that SNA96 did not address. We were unable to confirm any of SNA96's other observations. This is despite having a sample size over 4x larger than the original.

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. 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, our stimuli had no co-articulation between syllables, which may have made the stimuli seem less language-like. Third, 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. Note, however, that the distribution of scores appears to be roughly normal, and does not suggest a bipolar distribution with some attentive subjects and some inattentive subjects (see Fig. 1). We are currently conducting a follow-up replication to investigate these possibilities.

⁶An intermediate option is to transform accuracy measures to log-odds, though this method is less accurate than a logit model.

⁷Wald'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.

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.
- Bagwell, C., & Contributors, S. (2015). *Sox: Sound exchange, the swiss army knife of audio manipulation*. sox.sourceforge.net.
- 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.
- 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.
- IBM. (2017). *Watson developer cloud*. www.ibm.com/watson/developercloud/text-to-speech.html.
- 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.
- Saeften, 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.
- Team, A. (2017). *Audacity(r): A free, cross-platform digital audio editor*. <http://www.audacityteam.org/>.