Replication of Saffran, Johnson, Aslin, & Newport (1999) Statistical learning of tone sequences by human infants and adults, Exp. 1

Kara Murdock, Alyssa Savery, Kelsey Trimm, Juliani Vidal and Joshua K. Hartshorne

Department of Psychology, Boston College

Contact joshua.hartshorne@bc.edu

We replicated Exp. 1 of Saffran, Johnson, Newport, & Aslin (1999) Statistical learning of tone sequences by human infants and adults, as part of a multi-year effort to replicate every adult statistical word segmentation study. While we were able to replicate the finding of above-chance statistical segmentation of tone sequences, many of the other findings reported in the original paper did not replicate.

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, Aslin, and Newport (1999). This experiment was part of a larger project to systematically replicate as many experiments involving statistical word learning in adults as possible.

Method

Subjects

99 subjects were recruited through Amazon Mechanical Turk and paid for their participation. Subjects were randomly assigned to one of the two languages, with 50 subjects in Language One and 49 in Language Two.¹ We excluded five subjects who reported formal music activities (lessons, singing in a choir, etc.) in the last five years. We further excluded 17 who incorrectly answered any catch trials, leaving us with 77 subjects: 39 in Language One (mean age: 34, range: 21-56) and 38 in Language Two (mean age: 35, range: 24-65). All but one of these subjects were native speakers of English. All reported normal hearing.

Known differences with original.

- Our study had three times as many subjects (24 vs. 77) and thus considerably more power and precision.
- The original excluded subjects who had significant musical training since the 7th grade. Given the wide age range in our subjects, this seemed to risk confound. Thus, we used the 5-year criterion instead.

Possible differences with original.

- The age range for the original subjects is unknown.
- The native languages of subjects in the original study are unknown.

- It is unknown how and whether subjects in the original study were compensated (course credit, money, etc.).

Materials

Tone sequences were constructed using free online services; first it was necessary to pinpoint the exact frequencies of the tones in hertz using an Online Tone Generator (synzalski.com), and then the tones were created and downloaded using a Sine Tone Generator (audiocheck.net).

Tones were constructed out of eleven pure tones of the same octave (starting at middle C within a chromatic set) and were the same length (.33 s). In order to form six tone words, the pure tones were combined into groups of three. Language One consists of: ADB, DFE, GG#A, FCF#, D#ED, and CC#D. Some tones appeared in only one word (G# for example) while others occurred in multiple words (D for example). The six tone words were then used to create six different blocks that each contained 18 words. The tone words were concatenated in a random order without silence between words, and tone words did not occur twice in a row. Each word was concatenated together 70 times in order to produce a seven-minute continuous stream of tones, which was called Language One. A given tone stream, for example, would be DFEFCF#CC#DD#EDGG#A. As in the Saffran et al. study, there were no acoustic markers of word boundaries. Language Two was constructed in the same manner as the first, using the same original eleven tones but combined in different order to form six new words. Language Two consists of: AC#E, F#G#E, GCD#, C#BA, C#FD, G#BA.

Within both languages, transitional probabilities alone were the only consistent cues to the beginnings and ends of the tone words. In terms of Language One, the average transitional probability between tones within words was

¹We intended to recruit 50 for each, but ended up one short due to experimenter error.

0.64 (range: 0.42 - 1.00). In contrast, the average transitional probability between tones spanning word boundaries was 0.18 (range: 0.04 - 0.59). The statistical structure of Language Two was very similar to Language One. The transitional probabilities between the tones in Language Two averaged .71 (range 0.41 - 1.00), with lower average probabilities across word boundaries ($M=0.21$, range: 0.04 - 0.49).

A thirty-six item test was constructed in which participants were forced to select between two alternative choices in order to assess learning. Each test item was made up of two-tone sequences: a word and a non-word. Non-words consist of three-tone sequences that were also made using tones drawn from the language, but never occurred during exposure to the language (and thus had transitional probabilities of 0). One of the sequences presented on each trial was a word from Language One and the other was a word from Language Two. Thus, the memory set for subjects learning one language was the foil set for subjects learning the other. All six words from each of the two languages were paired for a total of 36 test trials. The tone sequences of each trial were separated by a 0.75 second interval of silence, and the trials themselves were separated by 1 second. We also added six catch trials that compared a real word against something that was obviously not in the training in order to find which subjects were not paying attention. They were always the last six items in the test.

Known differences with original.

- Saffran et al. (1999) generated stimuli using SoundEdit 16. We used audiocheck.net.
- Saffran et al. (1999) used a 5 second delay between each test trial. We used a shorter, 1 second delay because we did not want subjects to think the webpage was not loading.
- Saffran et al. (1999) generated two different random orders of the test trials. In contrast, we randomized the order of questions for each subject in order to help eliminate any effects of order.
- We randomized our training stimuli independently from Saffran et al., and this resulted in slight differences. Saffran et al. (1999) report an average between-word transitional probability of .14 (range: .05-.60) for Language One and .18 (range: .07-.53) for Language Two. In both cases, our average transitional probabilities were somewhat higher.
- Similarly, we find somewhat different within-word transitional probabilities. While the average for Language One was the same, the range was larger for them (0.25 - 1.00) than us (0.42 - 1.00). The same was true for Language Two, where their range was (0.33-1.00) and ours was (0.41 - 1.00). Note that this should make

it easier for subjects to learn our stimuli relative to theirs.²

Apparatus

Participants were recruited using Amazon Mechanical Turk through the Boston College Language Learning Lab and used their own computers and listening devices in order to listen to the tone streams and complete the experiment. The experiment was programmed using jsPsych (de Leeuw, 2014) and run on Amazon Mechanical Turk using PsiTurk (Gureckis et al., 2016).

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.

Procedure

Subjects were run independently and individually. Firstly, subjects were instructed that they would hear a recording of continuous tones. However, they were not told which aspects of the tone sequence would be tested. They were told to relax and try not to analyze the sound stream. Subjects then listened to the seven-minute-long recording of one of the two tone streams (Language One or Language Two) described above, repeated three times. Each of the three seven-minute listening sessions was followed by a short, subject-paced break. After a total of twenty-one minutes of listening, subjects were exposed to the 36-question test. They were instructed to indicate the most familiar tone sequence on each test trial. The correct choice for subjects exposed to Language One was the incorrect choice for subjects exposed to Language Two. Each test trial was randomized. All subjects heard three practice trials prior to testing.

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 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).
- In the original, subjects provided written responses.
- It is not clear how long the breaks between blocks were in the original, or whether these were subject-paced.

²It is not immediately clear whether this difference is within what one might expect based on random variation. We did not investigate this issue further.

Results

Original Analyses. In Language One, the mean score for subjects was 24.1 (SE=0.8) correct out of 36 total questions (67%), which was significantly above chance ($t(38)=7.3$, $p=9.7 \times 10^{-9}$). Performance was above chance for five of the six words ($ps < .01$). The original study similarly reported above-chance performance on five of six words, though our at-chance word (FCF#) differed from theirs (ABD). Language Two had a mean of 24.4 (SE=0.8) correct answers (68%), which was likewise above chance ($t(37)=7.5$, $p=6.1 \times 10^{-9}$). All six words in Language Two were learned at above-chance levels ($ps < .05$). Performance was not significantly different across the two languages ($t(75)=0.27$, $p=.79$).

After that, we analyzed how the transitional probabilities between tones might play a role in how well particular words were learned. Since each word contained three tones, there were two transitional probabilities associated with each word: between the first and second tones and between the second and third tones. For the following analysis, the average of the two transitional probabilities for each word was computed. We split the six words of each language into two sets, one set containing the three words with the highest average transitional probabilities (1.0, 0.75, and 0.75 for both languages), and the other set containing the three words with the lowest average transitional probabilities (0.420, 0.435, and 0.50 for Language One; 0.410, 0.665, and 0.665 for Language Two). A paired t-test showed no difference in subject accuracy for high vs. low transitional-probability words for Language One (69% vs. 65%; $t(38)=1.24$, $p=.22$) or Language Two (68% vs. 68%; $t(37)=0.13$, $p=.89$). This contrasts with Saffran et al. (1999), who reported significantly higher accuracy for the high transitional-probability words in each language. (Note that it is not clear how Saffran et al. conducted this analysis. They report using ANOVAs on subject means. However, the reported degrees of freedom are inconsistent with this.)

The original notes that the accuracy they observed for tones (Lang 1: 74%, Lang 2: 80%) was not significantly different from what was reported by Saffran et al. (1996) Exp. 1 (76%). The accuracy we observe is considerably lower (Lang 1: 67%, Lang 2: 68%), despite our stronger attentional filter. Nonetheless, it is very similar to what we observed in a recent replication of Saffran et al. (1996): 65%. Due to the relatively close matches, we did not perform statistical analyses.

Revised Analyses. There are two important limitations to the original analyses conducted by Saffran et al. (1999). 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).³

Additionally, they treat items as fixed effects. Thus, any

conclusions are specific to the 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 Saffran et al. (1999) did) and *also* 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 using binomial mixed effects models, implemented in R using the lme4 package, ver. 1.1-12 (Bates et al., 2017). We considered main effects of language and transitional probability (high/low), as well as their interaction. We set the contrasts to center the variables, such that the intercept is equal to the grand mean. We considered two possible random effects structures. The first estimated random intercepts for each word/foil pair and for each subject, as well as a random slope of transitional probability (high/low) for each subject. The second estimated separate random intercepts for each word and foil separately, random intercepts for each subject, and a random slope of transitional probability. The former provided a better fit as assessed by conditional AIC (7431 v. 7438). (We calculated cAIC using Stein formulas, as implemented in Greven & Kneib (2010) and Saefken & Ruegamer (2015)). Thus, we use the former random effects structure throughout, including random slopes where appropriate.

We found a significant intercept ($B = 0.82$, $SE=0.11$; Wald's $z = 7.8$, $p = 8.5 \times 10^{-15}$), indicating overall above-chance learning.⁴ The main effect of language was not significant ($B = 0.02$, $SE=.09$; Wald's $z = 0.20$, $p=.84$), nor was the main effect of transitional probability ($B = 0.07$, $SE=0.06$; Wald's $z = 1.15$, $p=.25$). The interaction was likewise not significant ($B = 0.09$, $SE=0.06$; Wald's $z = 1.51$, $p=.13$).

We followed these analyses with binomial mixed effects models for each trained word individually, with a random intercept of foil. The intercepts were significant for every word ($ps < .05$) with the exception of FCF#.

Thus, the mixed effects regression analyses supported the findings from the previous section based on t-tests.

Discussion

Our replication provided mixed support for the conclusions of Saffran et al. (1999).

As in the original, we found statistically significant learning of tone sequences, with evidence of successful learning for all but one of the trained words. Performance did not differ significantly across the two languages.

In contrast to the original, we found no evidence of superior learning for trained words with high transitional probabilities. This is consistent with our failure to observe such

³An intermediate option is to transform accuracy measures to log-odds, though this method is less accurate than a binomial model.

⁴Note that because of the contrast structure, these parameter estimates are half of the estimated effect size.

effects in our replications of Saffran et al. (1996), which used linguistic materials (e.g., Garcia et al., 2017). The most obvious explanation is that these findings are based on very small numbers of items, so failures to generalize to new items are perhaps to be expected.

Another place where our findings diverged was in learning success: We found much lower learning success than was reported by Saffran et al. (1999). This again matches what we saw for our replications of Saffran et al. (1996). Again, the most likely culprits are methodological. While Saffran et al. (1999) give little information about their subjects, if they were undergraduates at the University of Rochester, it is likely that our subjects were somewhat older and with lower IQs than theirs. Either could result in poorer performance on statistical word segmentation. Our study was also carried out over the Internet. We attempted to ensure that subjects were attending to the stimuli by rejecting those who did not pass the catch trials – something Saffran et al. (1999) did not do – but we cannot rule out differences in attention or motivation.

Data and Code

Data, analysis code, and experiment code are available at osf.io/79nk6.

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
- Saeften, B., & Rugeamer, D. (2015). *Conditional akaike information criterion for lme4*. cran.r-project.org/web/packages/cAIC4/.
- Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999, Feb). Statistical learning of tone sequences by human infants and adults. *Cognition*, 70(1), 27–52.
- 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.