Replication of Finn & Hudson Kam (2008) The curse of knowledge: First language knowledge impairs adult learners' use of novel statistics for word segmentation, Exp. 3

Gina L. Iozzo, Katie E. Lamirato, and Joshua K. Hartshorne

Department of Psychology, Boston College

Contact joshua.hartshorne@bc.edu

We replicated Exp. 3 of Finn & Hudson Kam (2008) The curse of knowledge: First language knowledge impairs adult learners' use of novel statistics for word segmentation, *Cognition*, 108, 477-499. 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 report a replication of Experiment 3 from Finn & Hudson Kam (2008), henceforth FHK08. This experiment was part of a larger project to systematically replicate as many experiments involving statistical word learning in adults as possible.

Method

Subjects

50 subjects were recruited through Amazon Mechanical Turk and paid for their participation. We excluded 11 subjects who incorrectly answered more than one catch trial, leaving 39 (Ages 23-61, $M=33$). All subjects were native speakers of English and reported normal hearing.

Known differences with original.

- Our study had nearly double the number of subjects (20 vs. 39) and thus considerably more power and precision.
- Subjects in our study were likely older than the undergraduate subjects in the original.
- Subjects in our study were compensated monetarily rather than with course credit.

Materials

Training. The stimuli cannot be easily recreated: The description in the original paper uses an obsolete phonetic alphabet for which we were unable to locate a key. The authors graciously provided us with the original stimuli. These stimuli are now available as part of the repository associated with the present project, with permission of the authors (osf.io/2v3ea/).

Stimuli consisted of eight two syllable words (CCVCV), each beginning with a consonant cluster. The CC onsets violate the word-initial phonotactic rules of English. FHK08 report that the stimuli were generated with the text-to-speech program SoftVoice. The synthesizer produced syllables with a monotonic F0 (fundamental frequency) of 83.62 Hz. All vowels were matched for length, and there were no co-articulation effects.

During training, words were presented quasi-randomly with no pauses and no immediate repetitions. FHK08 report transitional probabilities of 1.0 for syllable transitions that are word internal and .143 at word boundaries. They further report that phoneme transitional probabilities (PTPs) within words were higher than those across word boundaries, with word-internal PTPs, ranging from .25 to 1.0 and PTPs across the word boundaries ranging from .035 to .143. Thus, there was no overlap in the within-word and across-word PTPs.

Training lasted 17 min 59 sec. While FHK08 report that each word occurred 560 times, we suspect this is a type-o, since it would suggest 8.3 syllables/sec. The speech rate in the training files provided is much closer to 1 syllable/sec, which would be consistent with each word occurring 70 times for a *total* of 560 words in the training. However, we did not confirm this with an exact count of the words.

Test. After exposure, participants were given a forced-choice test between a trained word and either a non-word (consisting of two syllables from the language but a transitional probability of 0) and a split-word (a trained word minus the first consonant and with another trained word's initial consonant at the end, resulting in a CVCVC structure). In essence, the researchers constructed the split-cluster words by shifting one phoneme to the right in the exposure stimuli. This results in a foil that is a licit English word, pitted against a target which is not. The split-word foils measured the accuracy of parsing the consonant clusters in defiance of

English phonotactics, while the non-word foils served as a (high) baseline.

There were 8 of each type of test item (word vs. non-word, word vs. split-cluster word), yielding 16 test items in total. The two items in a pair were presented one after another with a 1 second pause in between. There was a 500 ms pause in between pairs during which participants were expected to answer. The order of the test items was randomized for each subject.

Upon inspection, one of the nonwords did not have an audible vowel in the first syllable, resulting in a CCCV structure.¹ In principle, this should have made it easy for the subjects to reject the nonword, though in fact, accuracy was only middling (0.62). In order to facilitate comparison with the original study, our primary analyses include this test item, though we also report analyses without that item.

In addition to these trials, we included eight catch trials that pitted a trained word against a nonword that included phonemes that never appeared in the training. These nonwords were created using MBROLA (Dutoit et al., 1996). Due to experimenter error, one of the catch trials did not have a correct answer, and so was excluded. The catch trials were always the last eight items in the test, and the order of the catch trials was randomized for each subject.

Known differences with original.

- FHK08 had a 3000 ms pause after test trials, whereas we used 500 ms. This was done to prevent subjects from thinking the web page was broken and refreshing.
- FHK08 did not include catch trials. Note that since our catch trials follow the critical trials, they cannot affect performance on the critical trials.
- FHK08 assigned subjects to one of two orders of test trials, each of which was the reverse of the other. We randomized ours independently for each subject.
- In FHK08, subjects responded to the test with paper and pencil. Our test was computerized.

Procedure

Participants were informed that they would be listening to a new language and were instructed to listen to it as best as they could, but were told to be careful not to overthink or ignore what they were listening to. After the more general instructions, participants were told, “*Kmodu* is a word in the language you are about to listen to.” This information was presented in written form, and the sound file for *kmodu* was also played. Before beginning training, the subjects were again informed that *kmodu* is a word in the language.

There were two separate listening blocks to allow subjects to take a break between the two sets of training. After this

exposure, they were told they would be tested to see what they learned about the language. They were asked to draw a picture during the exposure phase. After the experiment was over, we asked subjects to report what they drew. Only three reported failing to draw anything.

After exposure, participants were given the forced-choice test described above. They were instructed to listen to pairs of possible words and were asked to choose which word was a better example of the language. They were instructed to make their best guess if unsure. Participants indicated responses using two keys on their keyboard, pressing ‘1’ if the first item in the pair sounded closer to the language and ‘2’ if the second one sounded closer to the language. After the test, participants completed an additional survey consisting of questions regarding their demographic and language backgrounds.

Known differences with original.

- FHK08 do not report how many of their subjects drew pictures during the training. However, since they provided subjects with drawing materials, it is reasonable to assume that more did so in their study than in ours.
- Our subjects were tested online, whereas FHK08’s were tested in the laboratory.
- In the original, a live experimenter told each subject that *kmodu* was one of the words. In our replication, we presented that information visually and also played the sound file for *kmodu*.
- Subjects in our experiment used their own equipment, increasing variability across subjects. It is not clear what equipment was used by FHK08, but presumably it was more stable subject-to-subject.

Results

Original Analyses. Participants chose trained words over non-word foils more often than chance ($M=62\%$, $SE=8\%$; $t(38) = 3.6$, $p = .0009$). Participants also chose trained words over split-word foils at a rate exceeding chance ($M=58\%$, $SE=8\%$; $t(38)=2.2$, $p=.03$). These rates did not differ significantly ($t(38)=1.0$, $p=.33$). This contrasts with FHK08, where accuracy in the split-word foil condition was at chance. (They do not report the comparison between foil types).

FHK08 report ceiling accuracy (100%) on trials involving *kmodu* itself. We find a somewhat lower rate of 83% ($SE=4\%$).

This pattern of results was qualitatively identical when removing the bad test item mentioned above, removing the three subjects who failed to draw, or both.

¹In the repository, this is expnon5

Revised Analyses

There are two important limitations to the original analyses conducted by FHK08. 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 stimuli used. No assessment of generalization can be made. Statisticians have suggested two options. The first is to conduct analyses separately over subject averages *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 binomial mixed effects models, implemented in R using the lme4 package, ver. 1.1-12 (Bates et al., 2017). We included fixed effects for foil type (non-word vs. split-word) and random intercepts both for subject and for target/foil pair, plus a random slope of test type (nonword vs. split-word) by subject. We set the contrast structure to center the fixed effects, which results in the intercept being the grand mean.

We find a significant intercept, reflecting overall successful learning of the trained words ($\beta = .64$, Wald's $z = 2.65$, $p = .008$). The main effect of foil type was not significant ($\beta = -0.09$, Wald's $z = 0.38$, $p = .70$). Removal of the bad item, the non-drawers, or both had minimal effect on the p-value of the intercept ($p = .01, .006, .009$, respectively). In all cases, the main effect of foil type was not significant ($ps > .6$).

We also ran models with each foil type separately. In no case did we see significant evidence of learning in the split-word trials ($ps > .1$). In the non-word trials, we find a significant intercept, reflecting evidence for learning ($\beta = .73$, Wald's $z = 2.30$, $p = .02$). Removal of the bad item, the non-drawers, or both had no effect on the p-value of the intercept ($p = .04, .02, .03$ respectively).

Discussion

FHK08 report that the results of their Experiment 3 confirm the results from their Experiment 1, in which subjects were able to distinguish trained words from non-words (which had one phonemic transition probability of 0) but not from split-word foils (which had non-zero phonemic transition probabilities that were nonetheless lower than those of trained words but followed the phonotactics of English).

Using the same analyses employed by FHK08, we did not confirm that finding. Rather, we found successful discrimination in both conditions and no difference between conditions. Results of our revised analyses are mixed. We did replicate the difference in significance (significant learning for nonword trials but not split-word trials). However, the

difference was significant (the learning success for the two foil types was statistically indistinguishable). While FHK08 do not report the comparison between conditions, a significant difference is critical to their conclusions.

Thus, we were unable to confirm the findings of FHK08. As with any replication, our replication is imperfect, differing from the original in subject population (undergraduates vs. AMT), venue (online vs. in lab), and in certain aspects of the procedure. Any of these could have affected the results.

That said, it is highly likely that low statistical power explains at least some of the differences in findings. The original study had few subjects (20) and even fewer items (16) for a total of 320 binary data points. While we increased the number of subjects, the number of items remains the same. Inspection of the data reveals considerable variability in accuracy for the different items, with a substantial number being below chance. This variability is no doubt why the mixed effects models were far less confident about rejecting the null hypothesis of no learning than were the t-tests.

Data and Code

All data and code are available through the Open Science Framework at osf.io/2v3ea.

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.
- Dutoit, T., Pagel, V., Pierret, N., Bataille, F., & Van der Vrecken, O. (1996). The mbrola project: Towards a set of high quality speech synthesizers free of use for non commercial purposes. In *Spoken language, 1996. icslp 96. proceedings., fourth international conference on* (Vol. 3, pp. 1393–1396).
- Finn, A. S., & Hudson Kam, C. L. (2008). The curse of knowledge: First language knowledge impairs adult learners' use of novel statistics for word segmentation. *Cognition*, 108(2), 477–499.
- 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

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