Replication of Frank, Goldwater, Griffiths, & Tenenbaum (2010): Modeling human performance in statistical word segmentation, Experiment 1

Jesse Mu, James Ledoux, Joshua K. Hartshorne

{mu.j, ledoux.ja, hartshoj}@bc.edu

Department of Psychology, Boston College

Abstract

We replicated Exp. 1 of Frank, Goldwater, Griffiths, & Tenenbaum (2010): Modeling human performance in statistical word segmentation, *Cognition*, 117(2), 107–125, as part of a multi-year project to replicate every published adult statistical word segmentation study. With some minor differences likely associated with a higher sample size, we largely replicate the main conclusion that longer sentence length increases word segmentation difficulty.

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

Introduction

In this paper, we replicated Experiment 1 of a recent statistical word-segmentation paper by Frank, Goldwater, Griffiths, & Tenenbaum (2010), henceforth FGGT10. This experiment is part of a larger, multi-year project to replicate as many published statistical word-segmentation papers as possible, using larger sample sizes. Note we do not replicate the computational modeling in the paper; rather, we are interested solely in the reproducibility of the sentence length effects described in Exp. 1.

Methods

Subjects

We recruited 400 subjects from Amazon Mechanical Turk, randomly allocating 50 subjects to each of eight sentence length conditions. Within each sentence length condition, each subject was randomly assigned a unique language out of 50 total languages. In other words, 50 languages were reused across sentence length conditions, but no language was used twice within a sentence length condition.

For analysis, we dropped 6 subjects whose native language was not English and 109 participants who missed more than one out of six catch trials (see Materials), leaving 285 participants, which were still reasonably evenly distributed across languages and sentence length conditions (32–39 subjects/condition).¹ The ages of the subjects ranged from 20 to 73 ($\mu = 35.7$), which is most likely older and more varied than FGGT10's predominantly undergraduate population, although they did not report ages.

¹We ran the analyses in this paper while keeping all 400 participants and while keeping only the 189 participants who missed no catch trials, but aside from some slightly different values observed and significance levels (e.g. $p < 0.01$ vs $p < 0.001$), we observed no qualitative differences in the results presented here.

Known differences with original

- Our study has almost 3 times as many subjects used for analysis (285 vs. 98), and thus greater statistical power.
- Subjects in our study were most likely significantly older than those in FGGT10.
- FGGT10 does not specify whether students received monetary compensation or course credit. Ours received monetary compensation.
- FGGT10 discarded subjects with performance greater than 2 standard deviations below the mean. We use catch trials to achieve a similar effect.
- We excluded far more subjects for low accuracy (109) than did FGGT10 (2). This could reflect a higher rate of non-compliant subjects in our study (perhaps because the study was run online) or it could reflect a much tighter filter (By definition, nearly all subjects will be no more than 2 standard deviations below the mean).

Materials

Each participant heard a unique and randomly generated sample from one of 50 randomly generated artificial languages. The lexicon of this language was generated by concatenating 18 syllables (/ba/, /bi/, /da/, /du/, /ti/, /tu/, /ka/, /ki/, /la/, /lu/, /gi/, /gu/, /pa/, /pi/, /va/, /vu/, /zi/, /zu/) into six words: two with two syllables, two with three syllables, and two with four syllables. All speech was generated using the MBROLA speech synthesizer (Dutoit, Pagel, Pierret, Bataille, & Vrecken, 1996) with the us2 diphone database. Following FGGT10, all consonants and vowels were 25 and 225ms in duration, respectively, and the fundamental frequency of the speech was 100 Hz.

Train Subjects were randomly assigned to one of eight sentence length conditions (1, 2, 3, 4, 6, 8, 12, or 24 words per sentence). As training data, subjects heard 600 words, 100 each of the 6 words from the lexicon, broken into sentences of the given sentence length. Sentences were generated by randomly concatenating words with no adjacent-word repetition. No breaks were present between words in the sentences (equal co-articulation between every syllable), but there was a 500ms break between sentences during training. Length of training data ranged from 7.7 (length-24) to 12.5 minutes (length-1); shorter sentence length conditions had longer training materials due to more sentences and thus more breaks.

Test Test materials consisted of 36 target-distractor pairs from the language. 30 of the pairs consisted of a word from the language paired with a “part-word” distractor. The part-word contained the same number of syllables as the real word it was paired with, and was composed of the end of one word and the beginning of another from the language. For each language, 5 part-word foils were randomly generated for each of the 3 word lengths (2, 3, and 4 syllables), resulting in 15 part-words. Except for the length-1 condition, part-word sequences appeared in the sentences of the corpus, albeit at lower frequencies than true words.

Each part-word was then paired exhaustively with real words of the same length, making 30 target-distractor pairs, each real word appearing 5 times with same-length part-words, and correspondingly, each part word appearing twice with same-length real words.²

To ensure subjects were paying attention to the task, we created 3 catch words, composed of syllables that never appeared in the training data (/ˈnɪ.mɪ/, /ˈhɒʊ.sɪ.jə/, /ˈtɛɪ.mou.nou.lə/). Like the part-words, each catch word appeared twice with each real word of the same length, resulting in 6 catch trials and a total of 36 test pairs.

Test pairs were the same for a specific language regardless of sentence length condition. Test pairs were shuffled so that no real or part-word occurred twice in a row, and specific pairs shuffled such that word presentation order was random. Each pair was played with a 500ms break between words during testing.

Apparatus

The experiment was conducted on Amazon Mechanical Turk using psiTurk (Gureckis et al., 2016) and jsPsych (De Leeuw, 2015). The experiment was carried out in subjects’ homes using their own computers. While we could not control the listening environment, subjects were asked to relocate to a quiet place, use headphones, and keep the computer sound sufficiently high.

Procedure

Subjects were instructed to listen to a nonsense language for 15 minutes, and told that they would be tested on how well they learned the words of the language. After listening, they were instructed to make two-option forced-choice decisions between the test pairs indicating which sounded more like a word from the language by pressing either “1” or “2” on the keyboard.

Known differences with original

- Languages in our study were unique to each subject within a sentence length condition but reused across sentence lengths (50 languages), while languages from FGGT10 were unique to each subject regardless of language condition (98 languages). We chose to reuse languages in order

²Note that there is some intentional repetition of part-words during testing, so that subjects do not easily notice repetition of the real words only.

to more tightly control the sentence-length manipulation, providing greater power to detect any effects. (Note that we account for this re-use of languages in our mixed effects model; see Results).

- Our study used the MBROLA us2 diphone database, while FGGT10 used us3. This was due to an (at the time of experiment design) unresolved bug in the us3 database.³
- FGGT10 does not precisely specify how to create part-words; we replicated the constraints mentioned to the best of our ability with the procedure above. The original study does not specify how many part-words there are, if part-words are repeated during testing, and if so, how many times.

Results

Data and analysis scripts can be found at osf.io/dznjpj.

Subject performance across sentence length condition is depicted in Figure 1. Like FGGT10, we observed significant variation among individual performance in line with classical word segmentation studies (Saffran, Newport, & Aslin, 1996), but a systematic trend in mean participant performance across sentence lengths.

Figure 1: Violin (density) plots of mean participant performance in each sentence length condition. A line connects the means of each group, with shaded area representing $\pm$ standard error. The dotted line represents chance performance.

Following FGGT10, we analyzed test data with a multilevel (mixed-effect) logistic regression model (Gelman & Hill, 2006), with a continuous fixed effect of word length, a categorical fixed effect of sentence length, and a random effect of participant identity. Unlike FGGT10, we also added a random effect of the specific language used in the task, since in our case 50 languages were reused among the 285 participants. We fit an additional model with a fixed effect of the interaction of word length and sentence length, and

³bugs.debian.org/cgi-bin/bugreport.cgi?bug=788188

like FGGT10 found no significant differences in model fit ($\chi^2(7) = 13.3$; $p = 0.06$).⁴

Unlike FGGT10, we observed a negative effect of word length ($B = -0.11$; $p < 0.001$). Intuitively, this means that subjects had a harder time correctly answering questions with longer-length words. Furthermore, unlike FGGT10, we observed significant coefficient estimates for all sentence lengths (Table 1), indicating that performance in all sentence length conditions was significantly higher than chance.

Sentence length	B	z
1	2.23	13.84***
2	1.68	11.01***
3	1.20	7.6***
4	1.35	8.63***
6	1.19	7.97***
8	1.12	7.13***
12	1.06	7.12***
24	0.93	6.12***

Table 1: Coefficients (B) for each length condition of the mixed effects logistic regression model along with z scores. *** indicates $p < 0.001$.

However, there was a similar consistent downward trend in the coefficients of our model. While FGGT10 used the difference in significance of model coefficients between longer and shorter sentence length conditions as evidence for a sentence length effect, our universally significant coefficients do not allow this conclusion. Thus, to determine whether this downward trend was significant, we collapsed the responses of each subject into total scores (# correct out of 30) and measured the correlation between subject score and sentence length, observing a significant negative relationship (Kendall’s $\tau = -0.29$; $p < 0.001$).

Thus, while our study differs in FGGT10 by observing performance significantly above chance for all sentence length conditions, we replicate the downward trend of performance versus sentence length and thus **we come to the same conclusion**: that increasing sentence length negatively affects word segmentation performance. The differences in significance are likely due to the greater sample size and statistical power in our study; the actual means of subject performance in the sentence length conditions are qualitatively quite similar to FGGT10 Figure 1.

Revised analyses

A somewhat simpler way of answering the same research question is to treat sentence length as a continuous – rather than categorical – predictor. This has the added advantage of solving the convergence issue noted above (after centering – but not rescaling – both the word length and condition length predictors). Again, we found no evidence

⁴This finding must be somewhat tempered by the fact that the more complex model failed to fully converge.

for an interaction between word length and sentence length ($\chi^2(1) = 0$; $p = 0.98$), and so we do not include said interaction term in further analyses. Once again in contrast to FGGT10, we observed a negative effect of word length ($B = -0.11$; $SE = 0.031$; $p < 0.001$; Cohen’s $d = -0.048$). Crucially, we observe a significant negative effect of sentence length ($B = -0.04$; $SE = 0.006$; $p < 0.001$; Cohen’s $d = -0.147$). Overall, success is much higher than chance (Cohen’s $d = 0.569$).

For purposes of comparison, we re-analyzed the original data from FGGT10 in the same way. Again, we found no evidence for an interaction between word length and sentence length ($\chi^2(1) = 0$; $p = 0.96$), and so we do not include said interaction term in further analyses. Now, we observe a significant negative effect of sentence length ($B = -0.06$; $SE = 0.012$; $p = 0.005$; Cohen’s $d = -0.248$). In contrast, the effect of word length continues to not be significant ($B = 0.001$; $SE = 0.028$; $p = 0.96$; *Cohen’s d* = 0.001).

Discussion

We confirm FGGT10’s main conclusion in Exp. 1 that word segmentation becomes considerably more difficult as sentence length increases. Unlike FGGT10, we find above-chance performance even for the longest sentence lengths. This difference could be due to our greater power or our stricter attention screen. Some evidence for this possibility comes from our re-analysis of FGGT10’s data, in which we find a numerically larger effect of sentence length than in our own data.

The other minor difference in our results is that we observed a significant negative effect of test pair word length on segmentation ability. While FGGT10 also observed a negative effect, it was not significant. This difference, too, can likely be attributed to the greater power in our study.

All code and data used in this replication – including the R Markdown source files for this write-up – are available at osf.io/dznnpj.

References

- De Leeuw, J. R. (2015). JsPsych: A javascript library for creating behavioral experiments in a web browser. *Behavior Research Methods*, 47(1), 1–12.
- Dutoit, T., Pagel, V., Pierret, N., Bataille, F., & Vrecken, O. van der. (1996). The mbrola project: Towards a set of high quality speech synthesizers free of use for non commercial purposes. In *ICSLP*.
- Frank, M. C., Goldwater, S., Griffiths, T. L., & Tenenbaum, J. B. (2010). Modeling human performance in statistical word segmentation. *Cognition*, 117(2), 107–125.
- Gelman, A., & Hill, J. (2006). *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press.
- 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 be-

havioral experiments online. *Behavior Research Methods*, 48(3), 829–842.

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.