

A large-scale deconvolutional study of predictability and frequency effects in naturalistic reading

Are there separable effects of a word’s frequency vs. predictability in human sentence comprehension? Recent work in cognitive science implicates prediction as a major organizing principle in human and animal cognition [1, 20, 11], and psycholinguists have long studied the role of prediction in human sentence processing and its relation to other comprehension mechanisms [16, 13, 15, 24, 8, 17, 14, 4]. Some prominent theories of word recognition claim that ease of lexical access is modulated by the strength of a word’s representation in memory, independently of contextual factors that guide prediction [2]. Other theories hold that apparent effects of frequency are underlyingly effects of predictability [17, 14]. A number of studies using constructed stimuli that factorially manipulate frequency and predictability have found separable additive effects of each variable, supporting the position that frequency and predictability index distinguishable influences on lexical processing (see [23] for a review). However, such studies usually use cloze estimates of predictability, which are known to have difficulty differentiating degrees of low contextual probability [21, 22]. Furthermore, while constructed stimuli afford direct control over linguistic variables, results may be influenced by task-specific artifacts and should therefore be complemented by naturalistic studies [3, 9, 19].

This study explores the generalizability of these findings to typical sentence comprehension by searching for separable effects of frequency and predictability during naturalistic reading. Although naturalistic data address the aforementioned concerns about ecological validity, they have their own potential shortcomings that are addressed here in various ways. First, deconvolutional time series regression (DTSR) is used to address the possibility of temporally overlapping response profiles that violate the independence assumptions of linear regression and may therefore confound model interpretation and hypothesis testing [18]. Second, held-out evaluation is used to incorporate model validity directly into statistical tests, avoiding approaches (e.g. likelihood ratio testing) that implicitly evaluate on in-sample data [25]. Third, the natural collinearity between frequency and predictability [3] is addressed through large-scale data, specifically three large naturalistic reading time corpora: Natural Stories (self-paced reading) [6], Dundee (eye-tracking) [12], and UCL (eye-tracking) [5]. The corpora contain over one million data points in total generated by 243 human subjects. Failure to distinguish effects of frequency and predictability would therefore raise doubts about the existence of such a separation in naturalistic sentence comprehension.

Predictability and frequency are operationalized using 5-gram and unigram language models (respectively), each computed using KenLM [10] trained on the Gigaword 3 corpus [7]. Models fit ShiftedGamma impulse response functions [18] to these variables, as well as to the nuisance variable *word length*, along with (eye-tracking only) *saccade length* and an indicator variable for whether the previous word was fixated. Furthermore, to capture trends in the response at different timescales, models contain linear effects for the word’s index in the sentence (*sentence position*) and document (*trial*). Following [18], in addition to the intercept, models also fit a convolved intercept (*rate*) designed to capture effects of stimulus timing. The response used in all corpora was log fixation duration (go-past for eye-tracking).¹ Outlier filtering is performed in each corpus following the procedures described in [18]. Approximately half the data in each corpus is used for fitting, with the remaining half reserved for held-out evaluation. Models include by-subject random intercepts as well as by-subject slopes and impulse response parameters for each predictor.² Held-out hypothesis testing use a “diamond” ablative structure (*5-gram surprisal* vs. *unigram logprob*) via paired permutation test of the by-item losses on the evaluation set, pooling across all corpora.³ If predictability and frequency effects are additive, all four comparisons should be significant.

As shown in Table 1, this is not the case. While results show evidence that both frequency and predictability in isolation reliably index processing difficulty (both improve significantly over the baseline), they show no effect of frequency over predictability and thus do not support the existence of separable effects. They are instead consistent with either (1) an account of apparent frequency effects as epiphenomena of predictive processing [17, 14] or (2) a more circumscribed role for frequency effects than constructed experiments suggest. The discrepancy between constructed and naturalistic settings presents a puzzle for psycholinguists, and further investigation may shed new light on the interplay between prediction and memory in naturalistic sentence comprehension.

¹The overall pattern of significance does not change when first-pass durations are used.

²By-word random intercepts are not included because of their potential to subsume frequency effects.

³To account for different error variances across corpora, errors are rescaled by the joint standard deviation of the errors from the full and ablated models by corpus.

Comparison	p -value
5-gram only vs. baseline	0.0001***
Unigram only vs. baseline	0.0001***
5-gram + Unigram vs. Unigram-only	0.0001***
5-gram + Unigram vs. 5-gram-only	0.1440

Table 1: Pooled hypothesis testing results.

Corpus	Effect estimate (log-ms)							
	SentPos	Trial	Rate	WordLen	SacLen	PrevFix	Unigram	5-gram
Natural Stories	0.0098	-0.0216	-0.3069	—	—	0.0158	-0.0018	0.0174
Dundee	-0.0085	-0.0052	-0.0277	0.0068	-0.0021	-0.0178	-0.0067	0.0117
UCL		0.0524	-0.1330	0.0023	0.0221	0.0778	0.0005	0.0184

Table 2: Effect estimates in log-ms by corpus, computed as the integral of the impulse response over the longest time offset seen in training [18]. Following psycholinguistic convention, unigrams were encoded as log-probabilities while 5-grams were encoded as negative log probabilities (surprisal), resulting in opposite signs. Only one estimate for *sentence position/trial* is reported for UCL because sentences were shuffled, rendering *sentence position* and *trial* identical.

References

- [1] Andreja Bubic, D Yves Von Cramon, and Ricarda I Schubotz. Prediction, cognition and the brain. *Frontiers in human neuroscience*, 4:25, 2010.
- [2] Max Coltheart, Kathleen Rastle, Conrad Perry, Robyn Langdon, and Johannes Ziegler. DRC: a dual route cascaded model of visual word recognition and reading aloud. *Psychological review*, 108(1):204, 2001.
- [3] Vera Demberg and Frank Keller. Data from eye-tracking corpora as evidence for theories of syntactic processing complexity. *Cognition*, 109(2):193–210, 2008.
- [4] Stefan L Frank and Rens Bod. Insensitivity of the Human Sentence-Processing System to Hierarchical Structure. *Psychological Science*, 22(6):829–834, 6 2011.
- [5] Stefan L Frank, Irene Fernandez Monsalve, Robin L Thompson, and Gabriella Vigliocco. Reading time data for evaluating broad-coverage models of English sentence processing. *Behavior Research Methods*, 45(4):1182–1190, 2013.
- [6] Richard Futrell, Edward Gibson, Harry J . Tily, Idan Blank, Anastasia Vishnevetsky, Steven Piantadosi, and Evelina Fedorenko. The Natural Stories Corpus. In Nicoletta Calzolari, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Koiti Hasida, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, Hlne Mazo, Asuncion Moreno, Jan Odijk, Stelios Piperidis, and Takenobu Tokunaga, editors, *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Paris, France, 5 2018. European Language Resources Association (ELRA).
- [7] David Graff and Christopher Cieri. English Gigaword LDC2003T05, 2003.
- [8] John Hale. A probabilistic earley parser as a psycholinguistic model. In *Proceedings of the Second meeting of the North American Chapter of the Association for Computational Linguistics on Language technologies*, pages 1–8, 2001.
- [9] Uri Hasson and Christopher J Honey. Future trends in Neuroimaging: Neural processes as expressed within real-life contexts. *NeuroImage*, 62(2):1272–1278, 2012.
- [10] Kenneth Heafield, Ivan Pouzyrevsky, Jonathan H Clark, and Philipp Koehn. Scalable modified Kneser-Ney language model estimation. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, pages 690–696, Sofia, Bulgaria, 8 2013.
- [11] Georg B Keller and Thomas D Mrcic-Flogel. Predictive Processing: A Canonical Cortical Computation. *Neuron*, 100(2):424–435, 2018.
- [12] Alan Kennedy, James Pynte, and Robin Hill. The Dundee corpus. In *Proceedings of the 12th European conference on eye movement*, 2003.
- [13] M Kutas and S A Hillyard. Brain potentials during reading reflect word expectancy and semantic association. *Nature*, 307(5947):161–163, 1984.
- [14] Roger Levy. Expectation-based syntactic comprehension. *Cognition*, 106(3):1126–1177, 2008.
- [15] Maryellen C MacDonald, Neal J Pearlmutter, and Mark S Seidenberg. The lexical nature of syntactic ambiguity resolution. *Psychological Review*, 101(4):676–703, 1994.
- [16] William D Marslen-Wilson. Sentence Perception as an Interactive Parallel Process. *Science*, 189(4198):226–228, 1975.
- [17] Dennis Norris. The Bayesian Reader: Explaining word recognition as an optimal Bayesian decision process. *Psychological review*, 113(2):327, 2006.
- [18] Cory Shain and William Schuler. Deconvolutional time series regression: A technique for modeling temporally diffuse effects. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018.
- [19] Cory Shain, Marten van Schijndel, and William Schuler. Deep syntactic annotations for broad-coverage psycholinguistic modeling. In *Workshop on Linguistic and Neuro-Cognitive Resources (LREC 2018)*, 2018.
- [20] Yosef Singer, Yayoi Teramoto, Ben D B Willmore, Jan W H Schnupp, Andrew J King, and Nicol S Harper. Sensory cortex is optimized for prediction of future input. *eLife*, 7:e31557, 2018.
- [21] Nathaniel J Smith and Roger Levy. Cloze but no cigar: The complex relationship between cloze, corpus, and subjective probabilities in language processing. In *Proceedings of the 33rd CogSci Conference*, 2011.
- [22] Nathaniel J Smith and Roger Levy. The effect of word predictability on reading time is logarithmic. *Cognition*, 128:302–319, 2013.
- [23] Adrian Staub. The effect of lexical predictability on eye movements in reading: Critical review and theoretical interpretation. *Language and Linguistics Compass*, 9(8):311–327, 2015.
- [24] Michael K Tanenhaus, Michael J Spivey-Knowlton, Kathleen M Eberhard, and Julie C E Sedivy. Integration of visual and linguistic information in spoken language comprehension. *Science*, 268:1632–1634, 1995.
- [25] Samuel S Wilks. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *The Annals of Mathematical Statistics*, 9(1):60–62, 1938.