

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Cross-language structural priming in recurrent neural network language models

Permalink

<https://escholarship.org/uc/item/7832j4vp>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

ISSN

1069-7977

Author

Frank, Stefan

Publication Date

2021

Peer reviewed

Cross-language structural priming in recurrent neural network language models

Stefan L. Frank (s.frank@let.ru.nl)

Centre for Language Studies, Erasmusplein 1
6525 HT Nijmegen, The Netherlands

Keywords: structural priming; bilingualism; neural networks; surprisal; garden-path sentences

Introduction

Recurrent neural network (RNN) language models that are trained on large text corpora have shown a remarkable ability to capture properties of human syntactic processing (Linzen & Baroni, 2021). For example, the fact that these models display human-like structural priming effects (Prasad, Van Schijndel, & Linzen, 2019; van Schijndel & Linzen, 2018) suggests that they develop implicit syntactic representations that may not be unlike those of the human language system. A rarely explored question is whether RNNs are also able to simulate aspects of human *multilingual* sentence processing (Frank, 2021) even though training RNNs on two or more languages simultaneously is technically unproblematic.

Tsoukala, Broersma, van den Bosch, and Frank (2021) developed a bilingual version of the Dual-Path RNN model of sentence production (Chang, Dell, & Bock, 2006). This model not only code switches between its two languages but can also simulate cross-linguistic structural priming (Khoe, Tsoukala, Kootstra, & Frank, 2020), which is well known to occur in bilinguals (Hartsuiker, Pickering, & Veltkamp, 2004, among many others). However, because the model can handle only hand-crafted, miniature languages, it remains an open question whether bilingual RNNs trained on large amounts of natural language emulate properties of human bilingualism. Here, we demonstrate that RNNs trained on English and Dutch sentences account for a particular garden-path effect and are sensitive to structural priming, both within and between languages.

Method

Materials Sentence (1) contains a local, structural ambiguity: “the princess” can be read as part of an object NP conjunction (“the king and the princess”) or as the subject of a new clause (“the princess protects”). This is known as the NP/S-coordination ambiguity. The latter reading turns out to be correct when the disambiguating verb (“protects”) appears, which leads to increased reading time (RT) in sentences like (1) compared to unambiguous variants where a comma is inserted before the conjunction (“...the king, and the princess”; Frank & Hoeks, 2019; Hoeks, Hendriks, Vonk, Brown, & Hagoort, 2006). Our simulations used 120 of such sentences in

Dutch (taken from Hoeks et al., 2006) and their structurally identical English translations.

(1) The wizard guards the king and the princess protects the prince with her life.

RNN models Five Long Short-Term Memory RNN models (Hochreiter & Schmidhuber, 1997), differing only in random initial connection weights, were trained on next-word prediction. The networks had a 300-unit input embedding layer, a single, 600-unit recurrent layer, and a 300-unit layer before the output layer. The training corpus consisted of nearly 17M sentences (225M word tokens) scraped from web sources (Schäfer, 2015), approximately equally divided between Dutch and English. These languages were randomly mixed at the sentence level, in a different order for each network training repetition. The joint lexicon comprised 36,648 word types, unmarked for language so that interlingual homographs received a single representation.

Procedure Word surprisal, which is well known to correlate with word RT (e.g., Smith & Levy, 2013), was recorded on the disambiguating verb. The simulated garden-path effect is the difference between the verb’s surprisal in an ambiguous (no comma) and an unambiguous (comma) target sentence. Successful structural priming reduces surprisal on a repeated structure, so a stronger garden-path effect should appear after an unambiguous compared to an ambiguous prime sentence.

All primes and targets were presented in both ambiguity conditions and in both languages. Each sentence item (presented up to and including the critical verb) formed a prime for all target sentences, excluding targets that were identical to the prime, or differed only in ambiguity or language. Following Chang et al. (2006) and van Schijndel and Linzen (2018), structural priming was modeled as error-based adaptation of connection weights, identical to network training. The learning rate for priming was set to 0.2. After completing each prime test, connection weights were reset to their pre-priming values.

Results

Figure 1 shows that a garden-path effect was indeed predicted after all four types of priming. This effect is stronger after an unambiguous than after an ambiguous prime, that is,




Figure 1: Mean garden-path effect after within- or between-language priming with an ambiguous or unambiguous prime sentence. Error bars indicate 95% CIs.

the model is sensitive to structural priming. Moreover, priming is effective both within and between language, although it is stronger within language. These conclusions were confirmed in a linear mixed-effect analysis with full by-network and by-item random effect structure: The priming effect was highly significant ($b = 0.031, z = 4.72, p < .0001$) and somewhat weaker between than within language ($b = -0.002, z = -4.96, p < .0001$). Qualitatively similar results were obtained from analyses on Dutch and English targets separately.

Discussion and Conclusion

After training on a large number of Dutch and English corpus sentences, RNNs displayed the garden-path effect caused by the NP/S-coordination ambiguity. More importantly, the networks were sensitive to structural priming, both within and between languages: The garden-path effect was stronger after priming with an unambiguous structure. Within-language priming was more effective than between-language priming, which is generally consistent with human sentence production experiments (Bernolet, Hartsuiker, & Pickering, 2013).

The current results further demonstrate that bilingual RNN language models capture realistic aspects of the human bilingual syntactic system. However, their predictions remain to be validated against data from human sentence reading experiments. To date, surprisingly few cross-language structural priming experiments included RT as the dependent variable,¹ although Weber and Indefrey (2009) did show RT effects of priming between German and English active/passive structures.

¹Or at least, surprisingly few of such studies have been published.

References

- Bernolet, S., Hartsuiker, R. J., & Pickering, M. J. (2013). From language-specific to shared syntactic representations: The influence of second language proficiency on syntactic sharing in bilinguals. *Cognition*, *127*, 287–306.
- Chang, F., Dell, G. S., & Bock, K. (2006). Becoming syntactic. *Psychological Review*, *113*, 234.
- Frank, S. L. (2021). Towards computational models of multilingual sentence processing. *Language Learning*, *71*, 193–218.
- Frank, S. L., & Hoeks, J. C. (2019). The interaction between structure and meaning in sentence comprehension: Recurrent neural networks and reading times. In *Proceedings of the 41st Annual Conference of the Cognitive Science Society* (pp. 112–118).
- Hartsuiker, R. J., Pickering, M. J., & Veltkamp, E. (2004). Is syntax separate or shared between languages? cross-linguistic syntactic priming in Spanish-English bilinguals. *Psychological Science*, *15*, 409–414.
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, *9*, 1735–1780.
- Hoeks, J., Hendriks, P., Vonk, W., Brown, C., & Hagoort, P. (2006). Processing the noun phrase versus sentence coordination ambiguity: Thematic information does not completely eliminate processing difficulty. *The Quarterly Journal of Experimental Psychology*, *59*, 1581–1599.
- Khoe, Y. H., Tsoukala, C., Kootstra, G. J., & Frank, S. L. (2020). Modeling cross-language structural priming in sentence production. In *Proceedings of the 18th International Conference on Cognitive Modeling*.
- Linzen, T., & Baroni, M. (2021). Syntactic structure from deep learning. *Annual Review of Linguistics*, *7*, 195–212.
- Prasad, G., Van Schijndel, M., & Linzen, T. (2019). Using priming to uncover the organization of syntactic representations in neural language models. In *Proceedings of the 23rd Conference on Computational Natural Language Learning*.
- Schäfer, R. (2015). Processing and querying large web corpora with the COW14 architecture. In *Proceedings of the 3rd Workshop on the Challenges in the Management of Large Corpora* (pp. 28–34). Institut für Deutsche Sprache.
- Smith, N. J., & Levy, R. (2013). The effect of word predictability on reading time is logarithmic. *Cognition*, *128*, 302–319.
- Tsoukala, C., Broersma, M., van den Bosch, A., & Frank, S. L. (2021). Simulating code-switching using a neural network model of bilingual sentence production. *Computational Brain & Behavior*, *4*, 87–100.
- van Schijndel, M., & Linzen, T. (2018). A neural model of adaptation in reading. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- Weber, K., & Indefrey, P. (2009). Syntactic priming in German–English bilinguals during sentence comprehension. *NeuroImage*, *46*, 1164–1172.