Recurrent neural networks use discourse context in human-like garden path alleviation

Forrest Davis and Marten van Schijndel

Cornell University

fd252@cornell.edu

CUNY Conference on Human Sentence Processing, March 2020

Background

The knight killed ...

(a) the dragon and fell to the ground.
(b) by the dragon fell to the ground.

- In *isolation*, (b) can lead to a reading time slow down, known as a **garden path effect**

Context alleviates garden-path effects

Humans can use *discourse context* to **accommodate** conversational implicatures of the relative clause in (b) (Altmann and Steedman, 1988; Spivey-Knowlton et al, 1993; Trueswell and Tanenhaus, 1991)

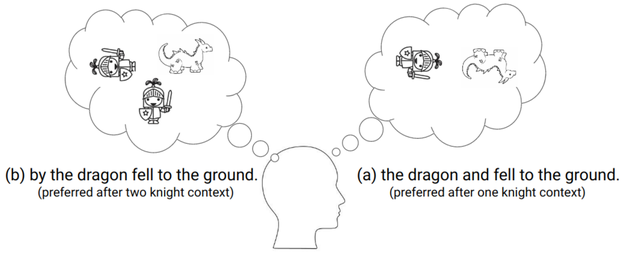

- Referent contexts (i.e. contexts with two referents) alleviate garden-path effects
- Temporal contexts (i.e. future contexts) alleviate garden path effects

Neural Network Language Models

- 10 RNN LMs trained on 80 million words of Wikipedia text
- 5 trained on data shuffled by sentence to remove discourse cues during training (as in the widely used model from Gulordava et al. 2018)
- 5 trained on sentences in their original order
- Measured information theoretical surprisal (Hale, 2001; Levy, 2008)

$$\text{surprisal}(w_i) = -\log \text{P}(w_i|w_1 ... w_{i-1})$$

Research Question

Humans have been claimed to use event knowledge for encoding discourse. Can the knowledge used in human-like garden path alleviation be learned from linguistic data alone, using recurrent neural network (**RNN**) language models (**LMs**)?

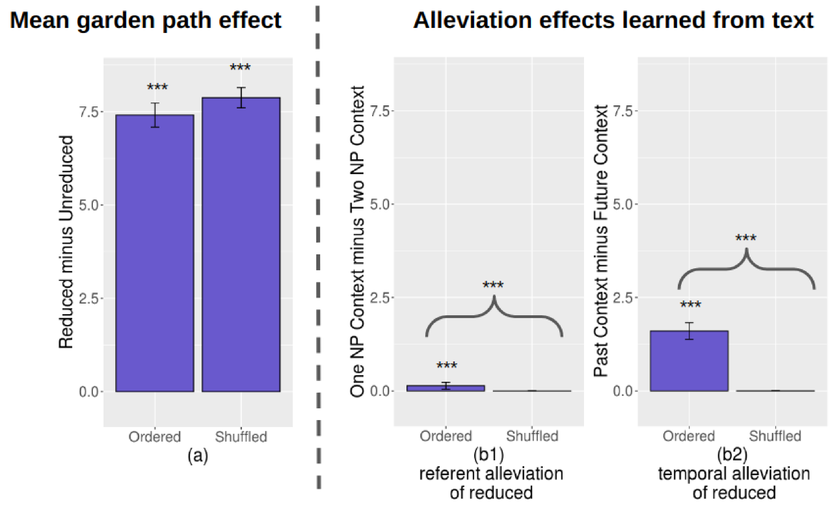

Figure 1: Differences between RNN LM surprisals. (a) between the verb+by region in a reduced versus unreduced relative clause (garden path effect). (b1) and (b2) when the reduced relative verb+by region is preceded by a garden-path supporting context ((b1) contexts with a single referent and (b2) past contexts) and when preceded by a garden-path alleviating context ((b1) contexts with two referents and (b2) future contexts).

Stimuli

a. **Context**
   (i) 1NP - A knight and his squire were attacking a dragon. With its breath of fire, the dragon killed the knight but not the squire.
   (ii) 2NP - Two knights were attacking a dragon. With its breath of fire, the dragon killed one of the knights but not the other.
b. **Target**
   (i) Reduced - The knight killed by the dragon fell to the ground with a thud.
   (ii) Unreduced - The knight who was killed by the dragon fell to the ground with a thud.

from Spivey-Knowlton et al (1993)

a. **Context**
   (i) Past - Several students were sitting together taking an exam in a large lecture hall earlier today. A proctor noticed one of the students cheating.
   (ii) Future - Several students will be sitting together taking an exam in a large lecture hall later today. A proctor will notice one of the students cheating.
b. **Target**
   (i) Reduced - The student spotted by the proctor received/will receive a warning.
   (ii) Unreduced - The student who was spotted by the proctor received/will receive a warning.

from Trueswell and Tanenhaus (1991)

Conclusion

- Both training conditions (i.e. with or without discourse contexts) exhibit garden path effects
- Referent and temporal contexts reduce the garden-path effect for only the RNN LMs trained on ordered data
- Multi-sentence discourse effects on sentence processing can be acquired from linguistic data alone

Discussion

- RNN LMs learn some pragmatic uses of relative clauses (i.e. to distinguish between entities of the same type) and tense (i.e. discourse about the future is likely to continue referring to the future)
- Humans may use event knowledge (Altmann and Ekves, 2019), but RNN LMs can capture these patterns; therefore, these patterns are, to some degree, learnable from language data alone
- Points to two alternatives: 1) garden path alleviation may not be evidence of the use of event knowledge in humans, or 2) aspects of pragmatic reasoning are learnable from just language data.

References

[1] Gerry Altmann and Zachary Ekves. Events as intersecting object histories: A new theory of event representation. *Psychological review*, 126(6):817, 2019.

[2] Gerry Altmann and Mark Steedman. Interaction with context during human sentence processing. *Cognition*, 30(3):191–238, 1988.

[3] Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. Colorless green recurrent networks dream hierarchically. In *Proceedings of NAACL*. 2018.

[4] John Hale. A probabilistic earley parser as a psycholinguistic model. In *Proceedings of NAACL*, pages 1–8. Association for Computational Linguistics, 2001.

[5] Roger Levy. Expectation-based syntactic comprehension. *Cognition*, 106(3):1126–1177, 2008.

[6] Michael J Spivey-Knowlton, John C Trueswell, and Michael K Tanenhaus. Context effects in syntactic ambiguity resolution: Discourse and semantic influences in parsing reduced relative clauses. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, 47(2):276, 1993.

[7] John C Trueswell and Michael K Tanenhaus. Tense, temporal context and syntactic ambiguity resolution. *Language and Cognitive Processes*, 6(4):303–338, 1991.