A Neural Model of Adaptation in Reading

Marten van Schijndel and Tal Linzen

Department of Cognitive Science, Johns Hopkins University

vansky@jhu.edu

Abstract

Humans rapidly adapt their lexical and syntactic expectations to match the statistics of the current linguistic context. We show that adding a simple adaptation mechanism to a neural language model improves our predictions of human reading times compared to a non-adaptive model.

Model

Initial model: LSTM LM trained on 90M words of English Wikipedia[3]
Adaptive model: Update parameters after each new sentence

Validation Data

Natural Stories Corpus Futrell et al. (2018)
2 genres in 10 texts:

- Documentary: 3 texts
- Fairy Tales: 7 texts

Self-paced reading data

- 181 Subjects

Linguistic Accuracy

	Perplexity		
Data	Initial	Adaptive	Δ%
Adapt across corpus	141.49	86.99	−38.5
Adapt within Documentary	99.33	73.20	−26.3
Adapt within Fairy Tales	160.05	86.47	−46.0

Psycholinguistic Accuracy

Predict reading times with model surprisal:

$$\text{surprisal}(w_i) = -\log \text{P}(w_i \mid w_1 \ldots w_{i-1}) \qquad (1)$$

Does adaptive surprisal predict reading times better than non-adaptive surprisal?

	$\hat{\beta}$	$\hat{\sigma}$	t
Without adaptive surprisal:			
Sentence position	0.55	0.53	1.03
Word length	**7.29**	**1.00**	**7.26**
Non-adaptive Surprisal	**6.64**	**0.68**	**9.79**
With adaptive surprisal:			
Sentence position	0.29	0.53	0.55
Word length	**6.42**	**1.00**	**6.40**
Non-adaptive Surprisal	-0.89	0.68	-1.31
Adaptive Surprisal	**8.45**	**0.63**	**13.42**

Psycholinguistic Plausibility of Model Adaptation Timeline

Ambiguous / Disambiguation

(1) The experienced soldiers warned about the dangers $\boxed{\text{conducted the midnight}}$ raid.

(2) The experienced soldiers who were warned about the dangers $\boxed{\text{conducted the midnight}}$ raid.

Unambiguous

$\text{RT}_{(1)} - \text{RT}_{(2)}$ and $\text{surprisal}_{(1)} - \text{surprisal}_{(2)}$ reveal the difficulty of disambiguation.

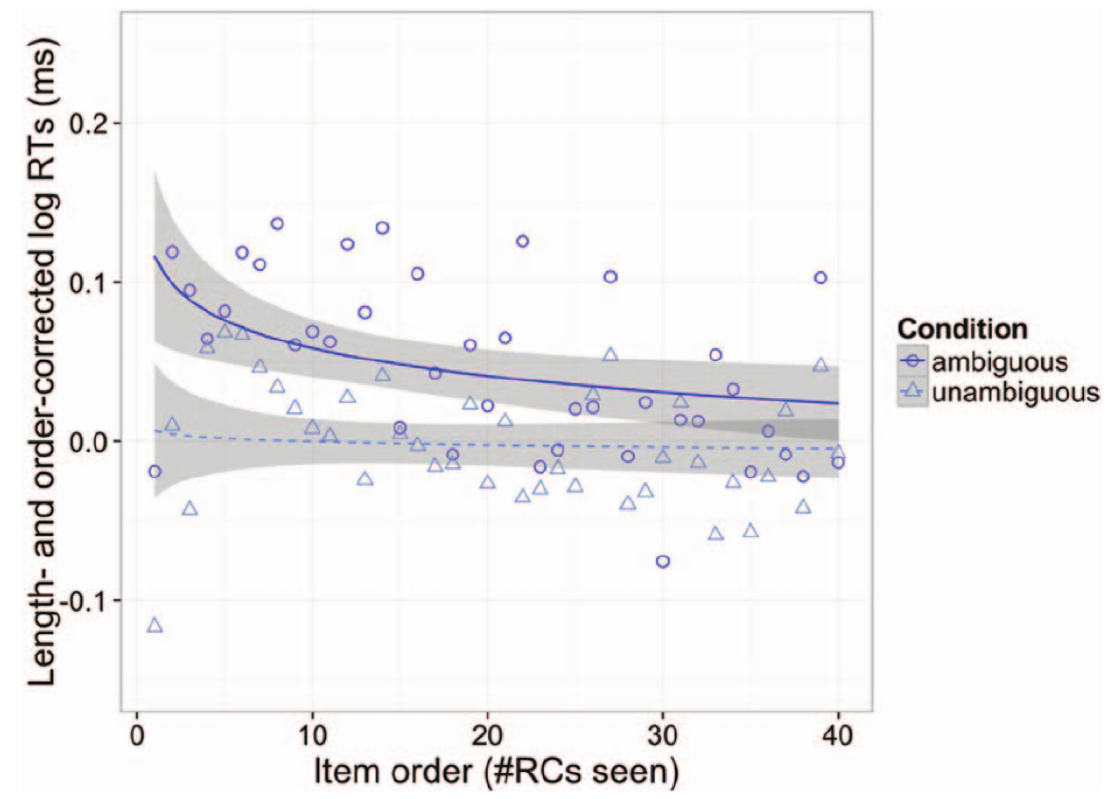

Figure 1: Human self-paced reading times (Figure from [1])

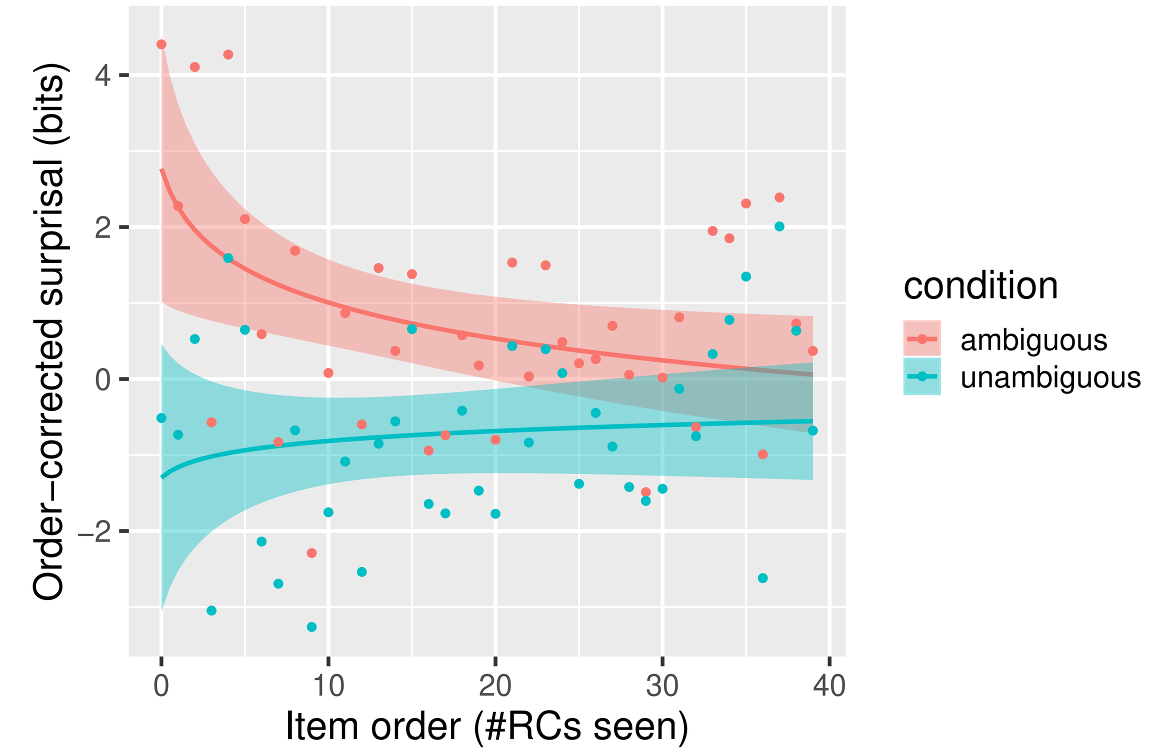

Figure 2: Adaptive model surprisal summed over the disambiguating region over the course of the experimental stimuli from [1]

Syntactic vs Lexical Adaptation

Use vocabulary-paired dative sentences to test how much adaptation is due to lexical experience versus syntactic experience.

1) Generate 200 dative sentences
   Prepositional object (PO): The boy threw the ball to the dog.
   Double object (DO): The boy threw the dog the ball.
2) Adapt to 100 DO items and 1000 Wikipedia sentences
3) Check perplexity on the held out 100 DO (shared syntax) and matched 100 PO (shared vocab) items
4) Repeat above 10 times for each of PO and DO adaptation (see paper for PO plots)
5) Repeat above for different learning rates

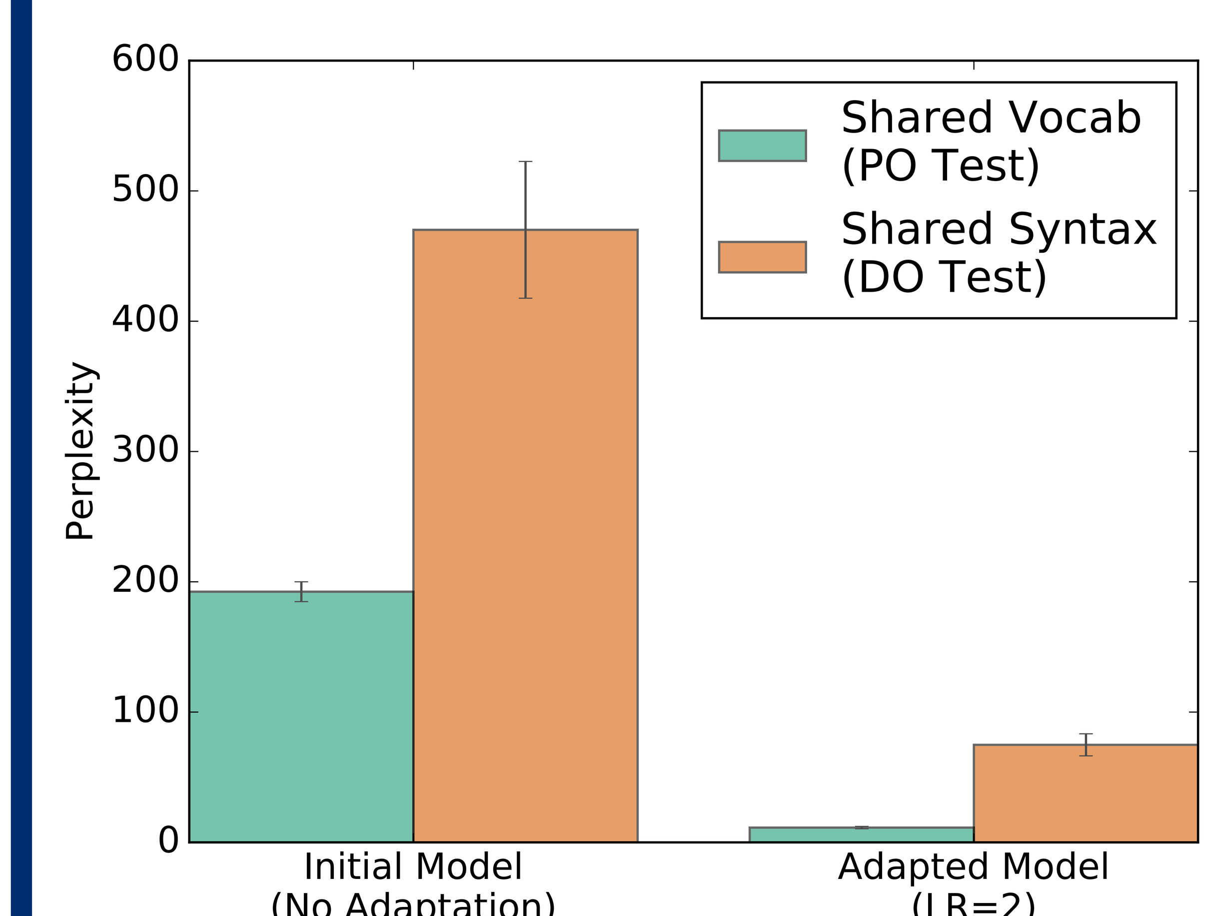

Figure 3: Adaptation to lexical and syntactic exposure.

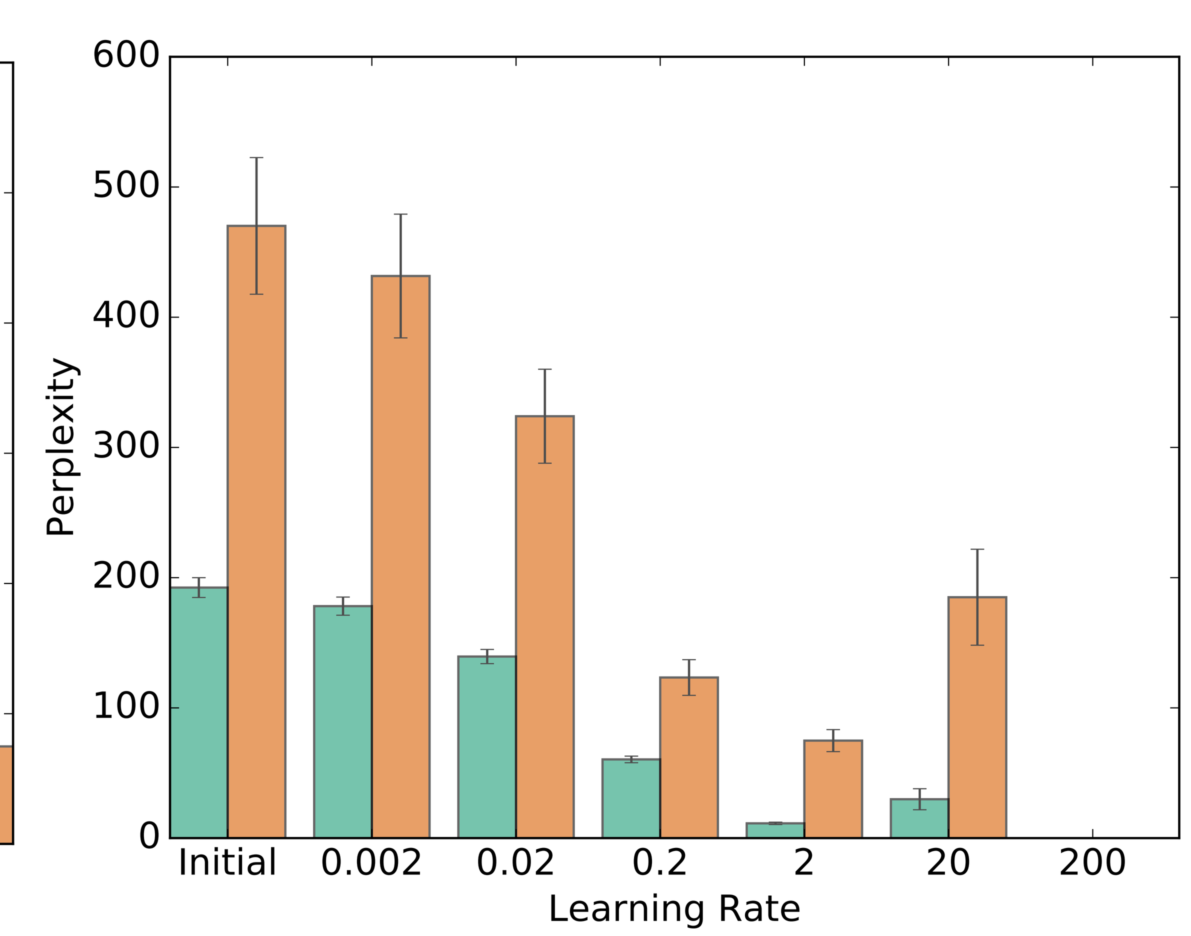

Figure 4: Learning rate influence over syntactic and lexical adaptation.

Conclusion

Model adaptation mimics psycholinguistic adaptation timeline observed by [1] by adapting to both syntax and vocabulary choice. Adaptation is very simple to implement and makes language models more linguistically and psycholinguistically accurate, and so should be adopted when using surprisal to model human cognition.

Beware Catastrophic Forgetting?

Multi-NLI corpus [4] has 10 genres of 2000 premise sentences each.

1) Adapt model to 1000 items from a genre ($G1$)
2) Adapt model to 1000 items from a different genre ($G2$)
3) Freeze weights and see if ($G1$) was unlearned

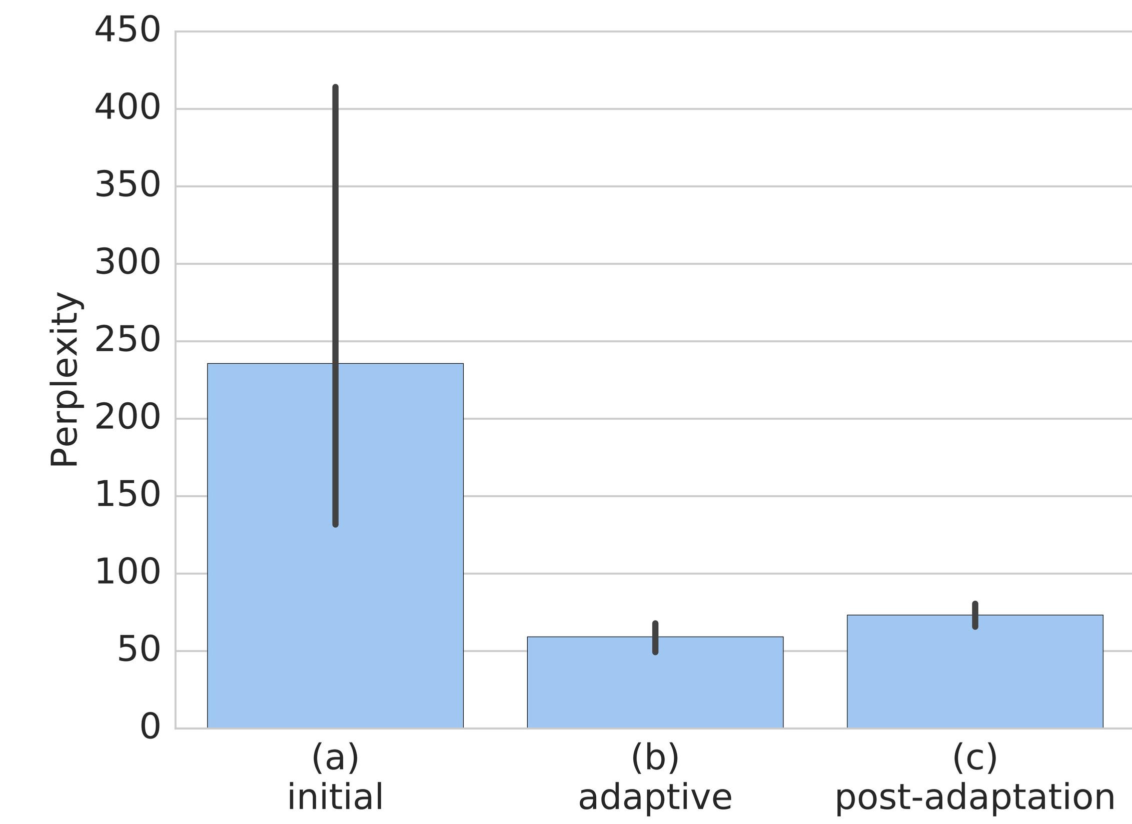

Figure 5: Perplexity on (a) $G1$ with no adaptation (b) $G1$ after adapting to $G1$ (c) $G1$ after adapting to $G1$ and $G2$

Catastrophic forgetting does not seem to be a problem with this amount of data.

References

[1] Alex B. Fine and T. Florian Jaeger. The role of verb repetition in cumulative structural priming in comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(9):1362–1376, 2016.

[2] Richard Futrell, Edward Gibson, Hal Tily, Anastasia Vishnevetsky, Steve Piantadosi, and Evelina Fedorenko. The natural stories corpus. In *Language Resources and Evaluation Conference*, pages 76–82, 2018.

[3] Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. Colorless green recurrent networks dream hierarchically. In *Proceedings of the North American Chapter of the Association for Computational Linguistics*, 2018.

[4] Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics, 2018.