

Supplementary Material

1 synthetic data

1.1 Number Agreement

For our experiments we used several 7 synthetically generated datasets, such that each contained sentences with one particular syntactic structure and varied lexical material. Each construction is stored in a distinct file. Each file contains 4 columns, which are separated by tabs:

1. The first column contains the sentence.
2. The second column contains the grammatical number of the subject (singular or plural).
3. The third column contains whether the main verb of the sentence agrees with the subject ('correct') or not ('wrong'). Each of the sentences is present twice, once with the correct agreement and once with the wrong one.
4. The fourth column contains the sentence id (which is the same for both correct and wrong version of the same sentence).

1.2 Tree Depth

For our regression experiments we used a large corpus with sentences with unambiguous but varied syntactic structures, generated by a script that follows a predefined context-free grammar. The output of this script is a four-column file containing:

1. The generated sentence.
2. The syntactic parse tree of the sentence, according to the grammar that was used to generate it.
3. The number of open nodes (syntactic tree-depth), following (Nelson et al, 2017).
4. The number of adjacent open and closing brackets before each word.

The columns of the file are separated with the character |.

We first processed a large corpus of such sentence with lengths between 2 and 25 words with our LSTM language model and stored the activations of all gates, the 2 hidden layers and memory cells of the model. Since syntactic depth is naturally correlated with word position, we filtered the processed words such that all position-depth combinations within positions 7-12 and depths 3-8 are uniformly represented in our final dataset. Note that the datapoints for our regression analysis are thus (word activation, tree depth) pairs. As a results from our sampling strategy, only one or a few pairs for each sentence in the original dataset are included. Our final dataset contains 4,033 positions from 1,303 sentences. We provide the file containing these sentences based on the above decorrelation method used in our study.

2 Evaluating model performance on the number-agreement task

To evaluate the performance of the LSTM language model on a given NA-task (Table 1), we tested for each sentence whether the model predicts the correct verb form (singular or plural) at the time point preceding the verb. Before presenting the model with a sentence, we first initialized its unit states by presenting a sequence of three arbitrary sentences from the test data¹. Then, the model was presented with all words up to one before the verb and the output probabilities of the model for the two verb forms were compared. If the probability of the correct verb form was higher then the wrong one then the trial was given a score of 1, else 0. Model performance on the NA-task was then defined as the average score across all sentences.

For the ablation studies, when a unit is ablated, we first set to zero all efferent and afferent weights to this unit before evaluating the model on the NA-task.

Nelson, M. J., El Karoui, I., Giber, K., Yang, X., Cohen, L., Koopman, H., ... & Dehaene, S. (2017). Neurophysiological dynamics of phrase-structure building during sentence processing. *Proceedings of the National Academy of Sciences*, 114(18), E3669-E3678.

¹In service, the aircraft was operated by a crew of five and could accommodate either 30 paratroopers, 32 unk and 28 sitting casualties, or 50 fully equipped troops. eos He even speculated that technical classes might some day be held for the better training of workmen in their several crafts and industries. eos After the War of the Holy League in 1537 against the Ottoman Empire, a truce between Venice and the Ottomans was created in 1539. eos Moore says: "Tony and I had a good unk and off-screen relationship, we are two very different people, but we did share a sense of humour". eos unk is also the basis for online games sold through licensed lotteries. eos"