

Lil-Bevo: Explorations of Strategies for Training Language Models in More Humanlike Ways

Venkata S Govindarajan✉ Juan Diego Rodriguez✉ Kaj Bostrom✉ Kyle Mahowald✉

✉Department of Linguistics ✉Department of Computer Science

The University of Texas at Austin

{venkatasg, juand-r, kaj, kyle}@utexas.edu

Abstract

We present Lil-Bevo, our submission to the BabyLM Challenge. We pretrained our masked language models with three ingredients: an initial pretraining with music data, training on shorter sequences before training on longer ones, and masking specific tokens to target some of the BLiMP subtasks. Overall, our baseline models performed above chance, but far below the performance levels of larger LLMs trained on more data. We found that training on short sequences performed better than training on longer sequences. Pretraining on music may help performance marginally, but, if so, the effect seems small. Our targeted Masked Language Modeling augmentation did not seem to improve model performance in general, but did seem to help on some of the specific BLiMP tasks that we were targeting (e.g., Negative Polarity Items). Training performant LLMs on small amounts of data is a difficult but potentially informative task. While some of our techniques showed some promise, more work is needed to explore whether they can improve performance more than the modest gains here. Our code and models are available online.¹

1 Introduction

Large Language Models (LLMs) generate complex and largely grammatical strings and display impressive performance with structures traditionally thought to require abstract and hierarchical syntax (Linzen et al., 2016; Linzen and Baroni, 2021; Wilcox et al., 2022; Futrell and Levy, 2019). They have achieved human-like performance at a wide range of natural language tasks (Bubeck et al., 2023; Frank, 2023), particularly those having to do with linguistic *form* (Mahowald et al., 2023). This state of affairs has led to claims that such models should be taken seriously as cognitive models of human language (Piantadosi, 2023; Baroni, 2022;

Frank, 2023), in line with claims from the neuroscience literature to “take mechanistic abstraction seriously” (Cao and Yamins, 2021).

One reason that has been posited *not* to take LLMs seriously as cognitive models, though, is the immense amount of data they are trained on relative to what a human child is exposed to (Warstadt and Bowman, 2022; van Schijndel et al., 2019). Thus, it is possible that models memorize more than humans do and, relative to humans, over-rely on statistical heuristics and memorized chunks of language (Bender et al., 2021).

On the other hand, the quality of data that LLMs get during pretraining is, in many ways, much worse than what human learners get. Children get richly structured, interactive, multimodal input, tailored to their specific interests and needs. A baby might reach for a cup of water and be told “Water. You want some water?” Given that babies are known to conduct repeated experiments to learn about the world (Gopnik et al., 1999), the baby might try this again and again until mastering the concept of what water is. An LLM, meanwhile, might begin learning language by being asked to predict random tokens in the Wikipedia article on quantum mechanics.

In this paper, we describe our experiments with Lil-Bevo, a small language model trained on human-scale data for the BabyLM competition (Warstadt et al., 2023). The goal of the competition is to train a performant LM on a human-scale amount of data: 10M words for the small track, 100M for the larger track. We submitted to both strict tracks — however, we were notified through the meta-review that our models qualify only for the loose track due to the usage of additional non-linguistic data (music from the MAESTRO dataset (Hawthorne et al., 2019)). The evaluation is on a set of natural language tasks including grammatical acceptability judgments via minimal pairs in the BLiMP benchmark (Warstadt et al.,

¹<https://github.com/venkatasg/Lil-Bevo>

2020a), language understanding tasks in SuperGLUE (Wang et al., 2019), and MSGS (the Mixed Signals Generalization Set) (Warstadt et al., 2020b)

We started with a baseline DeBERTa model, trained from scratch on BabyLM data using a custom unigram SentencePiece tokenizer (Kudo and Richardson, 2018). Our strategy was not focused on the architecture, but on ways in which we could adjust the training regime to improve performance above the baseline.

Specifically, our strategy targets 3 ways in which typical LLM training regimes lead to lower-quality data than humans have access to. Here, we describe those strategies and their motivation. We give detailed methods in Section 2 and then present results, including a number of ablation studies that attempt to partition out what strategies were successful.

We treated these studies as proof-of-concept and did not exhaustively test these strategies. Thus, we think that there is still room for improvement.

Training on Short Sequences Unlike LLMs, babies do not start language by learning long complicated sequences all at once. Using databases of child and child-directed speech, it has been shown that there is some alignment of caretakers to the child’s level in terms of linguistic complexity such that caregivers talk to younger children using shorter utterances and longer utterances as they develop (Schwab and Lew-Williams, 2016; Kunert et al., 2011). To that end, Mueller and Linzen (2023) showed that training on simpler data first could induce a better hierarchical bias for learning language. We specifically take inspiration from Press et al. (2021) who showed that LLMs learn better when trained on shorter sequences before being trained on longer sequences.

Training on Music Before Training on Language

Unlike LLMs, babies are exposed to a wide range of input besides just text. Before and while learning language, they are also learning to map the visual world, to navigate the physical world, to process non-linguistic auditory stimuli, and to engage in a wide variety of cognitive operations. Thus, it is commonly observed that some of the machinery thought to be language-specific (e.g., hierarchical structure) might be induced in pre-linguistic infants through exposure to other kinds of stimuli. Papadimitriou and Jurafsky (2020) use this idea to show that training language models on structured data (e.g., music) can help models learn faster. We

use a similar idea, with initial pretraining on a mix of music (piano performances) and text.

Targeted Masked Language Model The role of child-directed speech in human language learning is controversial (see Consortium and et. al., 2020, for discussion and a large-scale replication of infant-directed speech preferences). It is generally agreed that parents do not correct a child every time they make a grammatical error (Marcus, 1993), but there is also evidence that social feedback acts as a signal (Tomasello, 1992) and that parents structure input to be helpful (Weisleder and Fernald, 2013). When a child says something wrong, a parent might “recast” the utterance or highlight grammatical features that children are struggling with (Nicholas et al., 2001). Inspired by this idea, targeting the BLiMP (Warstadt et al., 2020a) syntactic evaluations as well as more general tasks, we trained with a targeted MLM objective.

We considered some variations of the idea of learning with some external feedback that distinguishes correct tokens against corrupted/noisy replacements. For example, ELECTRA (Clark et al., 2020) consists in learning to detect tokens which have been replaced by an auxiliary model. Unfortunately, replaced token detection approaches such as ELECTRA (Clark et al., 2020) suffer from an inability to learn probability distributions over the entire vocabulary, and so cannot be used for (pseudo)-likelihood scoring (Salazar et al., 2020). Another related approach is Corrective Language Modeling (CLM) (Bajaj et al., 2022), in which the model is trained to correctly replace corrupted tokens; however, it is not clear how to best use these models for scoring sentences in BLiMP.²

Given the problems outlined above, we decided to use masked language modeling (MLM) with targeted masks. The motivation is to make it easier for the model to learn syntactic phenomena that co-occur frequently with certain words. Other strategies for selecting masks were used in Sadeq et al. (2022); Gu et al. (2020); unlike these works, we mask specific words which are essential to the phenomena in BLiMP. For example, to target the filler-gap dependency subtask in BLiMP, we go through the original data set and mask every occurrence of “that” and “what” in the corpus. By

²Initial experiments with CLM performed worse than masked language modeling (MLM); we believe this is due to a mismatch between training and how the pseudo-likelihood scoring is done via masking.

focusing on these words, we anticipate that the model will more quickly learn to score “I know what you did last summer.” more highly than “I know that you did last summer.”

2 Experiments & Methods

We report all experiments and results for Lil-Bevo in this paper, as it enabled quick prototyping, and because we find similar trends with our larger model Lil-Bevo-X. Lil-Bevo-X differs from Lil-Bevo in the model used (deberta-base rather than deberta-small), training data (100M versus 10M), and vocabulary size. Final results for the Lil-Bevo-X are available on our [online repository](#).

Tokenizer We trained a unigram SentencePiece tokenizer (Kudo and Richardson, 2018) from scratch on the BabyLM data combined with the MAESTRO (Hawthorne et al., 2019) dataset (described in detail below) using the sentencepiece library. Specifically, we trained a tokenizer with a vocabulary size of 16,640 and 33,280 for Lil-Bevo and Lil-Bevo-X respectively. `<mask>` and `<cls>` were included as control symbols in the vocabulary, along with an end-of-sequence token (`</s>`), a pad token (`<pad>`) and an unknown token (`<unk>`).

Model We chose to use an encoder-based language model, specifically DeBERTa since (a) encoder-based language models are known to capture many syntactic and semantic features in language when pretrained on relatively modest amounts of data (Zhang et al., 2021), (b) there were a wide variety of off-the-shelf DeBERTa architectures available on HuggingFace for easy prototyping and use.

We trained the model in three phrases: (1) pretraining on a combination of music and text for 5 epochs with a sequence length of 64 tokens, (2) continuing pretraining on text for 50 epochs with a sequence length of 128 tokens, and (3) finally pretraining on text using targeted MLM for 2 epochs with a sequence length of 512 tokens. Each of these is described in more detail below.

1. Music Pretraining Papadimitriou and Jurafsky (2020) find that pretraining on languages other than the target language — including music and code — lead to lower perplexities on target language as compared to random distributions of tokens, or even Zipfian token distributions. Inspired by this idea, we explored whether supplementing

the 10M linguistic tokens with *non-linguistic* musical tokens from the MAESTRO dataset (Hawthorne et al., 2019) could lead to noticeable improvements in LM learning. The impetus behind pretraining on music is two-fold: (a) additional training data that nevertheless has structural biases that could help the model learn structural biases found in language (b) the model reaching a stable region in parameter space that enables it to learn desired linguistic properties much faster and/or better.

After several experiments, we found that pretraining on the combined *strict-small* and the entire MAESTRO dataset for 5 epochs provided the best results. We use V3.0.0 of the MAESTRO dataset, which contains 85M tokens using our custom trained tokenizer. The dataset consists of 200 hours of MIDI piano recordings, which we convert to text and tokenize with the shared unigram SentencePiece tokenizer. Our textual representation of MIDI consists of a chronological sequence of codes describing the channel and key of each note onset and release event (e.g. `c0n71` for ‘note on, channel 0, key 71’) delimited by spaces and optional codes for time between events (e.g. `t18` for 18 MIDI ticks). We chose a short sequence length of 64 tokens for pretraining inspired by the Shortformer, which we now explain in further detail.

2. Shortformer Press et al. (2021) introduce a few innovations to the training regime. In particular, we focused on their idea of training for shorter sequence lengths before moving onto longer ones. We used a similar training regime to (Press et al., 2021), where we started with a training sequence length of 128 for 50 epochs, before moving to a training sequence length of 512. We initially experimented with training on longer subsequence length for 150 epochs as in Press et al. (2021), but discovered lower evaluation results on most BLiMP categories (albeit with some improvements on some categories like Island Effects and Quantifiers). Results on BLiMP (Warstadt et al., 2020a) and SuperGLUE (Wang et al., 2019) saturated with as little as 2 epochs — we believe this is because of the much smaller size of the dataset as compared to that in (Press et al., 2021), leading to overfitting on the dataset.

3. Targeted MLM We specifically masked out words which were essential to some of the BLiMP subtasks. Some of these, such as quantifier and negation words, are also important to some of the

| Category | Total | Avg |
|-------------------|--------|-----|
| S-V agreement | 124197 | 4.3 |
| Animacy | 100206 | 3.5 |
| Quantifiers | 89926 | 3.1 |
| Modal verbs | 58604 | 2.0 |
| NPI licensing | 47484 | 1.6 |
| Filler gap | 34988 | 1.2 |
| D-N agreement | 28675 | 1.0 |
| Adverbs | 19332 | 0.7 |
| Anaphor agreement | 3659 | 0.1 |

Table 1: Total number of masks and average number of masks per sample for each targeted category (*S-V agreement* stands for subject-verb agreement, and *D-N agreement* stands for determiner-noun agreement).

SuperGLUE tasks (e.g., textual entailment.) For anaphor agreement, we masked the words “himself”, “herself”, “itself”, “themselves”. For NPI licensing the masked words included “not”, “often”, and “probably”³. The list of words which were masked in each category are shown in Table 3 in Appendix A. We used a sequence length of 512 tokens, and additionally masked other random tokens in order to mask a total of 15% of tokens per sample.

The total number of words masked for each category across the 10M train set are given in Table 1.

The *Animacy* class consists of animate nouns, and was used to target the minimal pairs in the *Argument Structure* category with animate/inanimate subjects (“Amanda was respected by some *waitresses*.” vs “Amanda was respected by some *picture*”). To obtain a list of animate nouns we used all the lemmas of (direct and indirect) hyponym synsets of *person.n.01* in WordNet.

In addition to targeting the BLiMP categories of S-V agreement, quantifiers, NPI licensing, filler gap, argument structure, DN- agreement and anaphor agreement, we also included some *modal verbs* (e.g., can, might, shall) and certain *adverbs* (e.g., never, maybe, always, perhaps), since these are important for textual entailment.

2.1 Ablations

We compare Lil-Bevo with ablations to explore how important our three strategies are for final performance. Specifically, we compare Lil-Bevo with

³Note that the masked words are not necessarily NPI items themselves, but rather that they are targets of single word substitutions in NPI items.

the following:

Long-only Train DeBERTa with a sequence length of 512 tokens for 57 epochs.

Short-only Train DeBERTa with a sequence length of 128 tokens for 57 epochs.

Short+target Train DeBERTa with a sequence length of 128 tokens for 55 epochs. Then train with targeted MLM for 2 epochs.

Music+short Train DeBERTa on music and text for 5 epochs with a sequence length of 64 tokens. Then continue training on text with a sequence length of 128 tokens for 52 epochs.

Music+short+long Train DeBERTa on music and text for 5 epochs with a sequence length of 64 tokens. Then continue training on text with a sequence length of 128 tokens for 50 epochs, followed by training with a sequence length of 512 tokens for 2 epochs.

Lil-Bevo (music+short±target) This is the same as *Music+short+long* except that the final stage of pretraining for 2 epochs uses targeted MLM.

Implementation We train all our models using the Trainer API, part of the huggingface python package. Models are trained using 4 Nvidia A40 GPUs, with the maximum possible batch size that was permissible with each experiment. Apart from setting initial learning rate to 6e-4, weight decay to 0.1 and a warmup ratio to 0.0001, we use default training arguments in the API (except for the final targeted MLM/long stage, where we used all default parameters). Models are evaluated on the validation split of the BabyLM dataset. We did not use the test split of the BabyLM data. We release all of the above pretrained models [online on the Huggingface Hub](#).

3 Results

Results for BLiMP, MSGS, SuperGLUE and the supplementary tasks are shown in Figure 1. The results are color-coded to represent each model’s differences from the Short-only ablation. We highlight some results below.

Does pretraining on music help? Comparing *short-only* with *music+short*, we see that pretraining on music helps slightly on 8 of the 12 BLiMP subtasks, and on two of the 5 supplement tasks. However, it suffers from a large gap of 9.1 points

| | | | | | | | | |
|---|------------------|----------------|------------|-----------|---------------|------------------|----------|------------|
| S-V Agr. | 65.4 | 83.9 | 82.2 | 78.2 | 83.8 | 83.9 | 84.8 | BLIMP |
| Quantifiers | 70.5 | 66.7 | 72.7 | 69.4 | 71.3 | 73.1 | 68.7 | |
| NPI Licensing | 55.9 | 77.2 | 61.0 | 54.9 | 65.1 | 63.7 | 78.5 | |
| Island Effects | 39.9 | 44.1 | 50.8 | 44.5 | 58.3 | 55.5 | 55.8 | |
| Irregular Forms | 87.4 | 88.5 | 87.0 | 84.2 | 87.0 | 85.3 | 85.3 | |
| Filler-Gap | 63.5 | 76.1 | 76.3 | 72.1 | 77.1 | 76.0 | 77.5 | |
| Ellipsis | 76.4 | 87.5 | 84.7 | 82.5 | 85.1 | 82.5 | 82.0 | |
| D-N Agr. | 90.8 | 90.5 | 89.8 | 88.9 | 90.9 | 90.8 | 91.7 | |
| Control/Raising | 67.9 | 69.8 | 69.8 | 68.9 | 72.1 | 71.9 | 70.0 | |
| Binding | 67.3 | 64.1 | 72.4 | 69.5 | 72.9 | 72.6 | 63.3 | |
| Arg. Structure | 67.1 | 73.7 | 71.7 | 69.8 | 71.1 | 69.9 | 72.5 | |
| Anaphor agreement | 81.5 | 91.2 | 92.3 | 91.7 | 90.0 | 89.6 | 90.9 | |
| syntactic-category-relative-position | -45.0 | 31.3 | 33.2 | 30.1 | 29.3 | 29.4 | 31.6 | MSGS |
| syntactic-category-lexical-content-the | 16.3 | 31.0 | 37.4 | 38.3 | 37.0 | 36.3 | 37.5 | |
| main-verb-relative-token-position | -79.4 | 47.5 | 58.0 | 47.7 | 49.6 | 54.8 | 40.7 | |
| main-verb-lexical-content-the | -99.3 | 67.2 | 40.6 | 42.1 | 42.2 | 54.8 | 36.7 | |
| lexical-content-the-position | 100.0 | 72.4 | 79.9 | 77.4 | 99.5 | 97.8 | 88.2 | |
| control-raising-relative-token-position | -77.7 | 35.1 | 33.8 | 33.6 | 35.3 | 34.6 | 34.0 | |
| control-raising-lexical-content-the | -28.3 | 44.2 | 40.1 | 43.8 | 44.0 | 44.1 | 44.0 | |
| WSC | 61.4 | 61.4 | 60.2 | 59.0 | 60.2 | 61.4 | 61.5 | SuperGLUE |
| SST-2 | 87.0 | 87.4 | 87.2 | 89.0 | 88.4 | 87.2 | 88.4 | |
| RTE | 61.6 | 55.6 | 50.5 | 48.5 | 44.4 | 48.5 | 46.5 | |
| QQP (F1) | 73.7 | 85.0 | 85.2 | 85.0 | 85.2 | 85.5 | 85.5 | |
| QNLI | 77.0 | 82.1 | 81.9 | 82.2 | 80.9 | 81.6 | 81.6 | |
| MultiRC | 61.4 | 63.1 | 64.2 | 63.6 | 65.5 | 64.8 | 66.0 | |
| MRPC (F1) | 79.2 | 80.0 | 80.8 | 81.9 | 81.1 | 80.9 | 82.2 | |
| MNLI-mm | 74.0 | 75.2 | 75.6 | 76.0 | 76.3 | 76.0 | 76.3 | |
| MNLI | 73.2 | 75.0 | 74.9 | 75.4 | 75.9 | 75.0 | 75.4 | |
| COLA | 25.8 | 71.6 | 72.3 | 71.7 | 74.1 | 73.9 | 73.7 | |
| BoolQ | 66.3 | 65.6 | 66.4 | 66.0 | 64.2 | 65.3 | 65.4 | |
| Turn taking | 53.2 | 65.7 | 68.9 | 68.2 | 68.2 | 67.9 | 68.2 | supplement |
| Subj Aux Inversion | 71.7 | 74.7 | 77.3 | 79.4 | 79.2 | 79.0 | 76.5 | |
| QA Congruence tricky | 32.1 | 51.5 | 52.1 | 43.0 | 43.0 | 46.7 | 57.0 | |
| QA Congruence easy | 31.3 | 79.7 | 76.6 | 65.6 | 73.4 | 73.4 | 82.8 | |
| Hypernym | 49.4 | 47.1 | 48.1 | 47.1 | 49.4 | 48.7 | 48.1 | |
| | RoBERTa-baseline | Short - target | Short only | Long only | Music - short | Music-short-long | Lil Bevo | |

Figure 1: Results for each model, for each task. The color reflects the difference in score between the given model and the RoBERTa baseline results released by the organizers of BabyLM.

on *QA Congruence tricky*. On SuperGLUE, *music+short* outperforms *short-only* on 6 of the 11 subtasks, and only slightly. Thus, we do not think there is strong evidence that pretraining on music improves over the short-only condition, in isolation.

Comparing *Lil-Bevo* (music+short+target) with *short+target*, we see that *Lil-Bevo* outperforms *short+target* on 69% of all tasks. Predicting score for each task in a mixed-effect linear regression with a fixed effect predictor for whether the model was *Lil-Bevo* or *short+target*, we found that *Lil-Bevo* was slightly better ($\beta = 1.3$, $\chi^2(1) = 4.11$, $p < .05$ by a likelihood ratio test). So, while music pretraining may help, the effect is small and inconsistent in our observed data.

What is the effect of targeted MLM? We compare *music+short+long* with *Lil-Bevo* (music+short+target) and *short-only* with *short+target* to ascertain whether targeted MLM helps over random masking. Targeted MLM does not systematically improve performance, except for two BLiMP tasks: NPI Licensing and Argument Structure. For NPI Licensing, *Lil-Bevo* outperforms *music+short+long* by 14.8 points, and *short+target* outperforms *short-only* by 16.2 points. We suspect that this difference could be meaningful since our Targeted MLM strategy specifically targets NPI terms that are substituted in BLiMP.

The effect of increasing sequence length When comparing *music+short* with *music-short-long*, and *short-only* with *long-only*, we find that pretraining with 512-token sequence lengths generally underperforms pretraining with 128-token sequence lengths. The difference between *short-only* and *long-only* conditions is quite large in fact. A linear mixed effect regression comparing the two using the same method as above found that performance was 1.8 points worse on average for the *long-only* method ($\beta = 1.8$, $\chi^2(1) = 14.2$, $p < .001$ by a likelihood ratio test). Thus, we believe pretraining with shorter sequences helps significantly compared to using longer sequences.

4 Discussion

Overall, we found that, for BabyLM’s, sequence length matters, music pretraining may help a little (but may be spurious), and targeted MLM training may help on specific tasks.

These results are far from exhaustive, and we

| Model | Dynabench score |
|------------------|-----------------|
| Lil-Bevo | 0.64 |
| Music-short-long | 0.64 |
| Music-short | 0.69 |
| Short-only | 0.63 |
| Short-target | 0.62 |
| Long-only | 0.61 |
| Lil-Bevo-X | 0.69 |

Table 2: Scores on Dynabench for different models.

see a number of areas for future improvement using these methods. To fully understand the role of initial pretraining on music, one could construct a series of synthetically-generated music datasets, with varying degrees of complexity. Would pretraining on music that is more “language-like” (Lerdahl, 1996) in some sense improve performance on downstream tasks? Perhaps there is a principled way to interpolate between music and language, using the same kind of data format (MIDI). At one end of the spectrum one would have MAESTRO, and at the other end, text that has been encoded as into MIDI events.

Related to the use of varying sequence lengths, future work could consider improvements in data preprocessing and batching; in particular, knowing the beginning and ending of coherent chunks of text (e.g., dialogues or documents) could help improve the model. Beyond this, Mueller and Linzen (2023) provide some evidence that curriculum learning approaches may be fruitful to improving low-resource language models.

Finally, a more thorough analysis is needed on when (and by how much) targeted MLM is able to boost model performance. Other strategies are also possible, such as combining targeted MLM with information-theoretic strategies for picking random masks (Sadeq et al., 2022). Beyond MLM, contrastive objectives could be used to encourage the model to score grammatical sentences more highly than ungrammatical sentences.

5 Conclusion

A big motivating question for training models on human-scale data is whether it is possible for models to attain linguistic competence without the massive amounts of data used to train the massive LLMs that dominate NLP leaderboards. If so, that would make it more plausible that we should

take LLMs seriously as cognitive models. So can BabyLMs learn like grown-up ones? While we find some hints of directions to pursue for making small language models learn more from less, we did not come close to matching LLM performance from larger amounts of data. Of course, that does not mean it is not possible to do so, and other teams might have different experiences. We did not fully explore optimizing all of our methods, and we treated our manipulations largely as proof-of-concept. Aggregating methods and results from a wider variety of teams will make it possible to more fully explore these questions.

References

- Payal Bajaj, Chenyan Xiong, Guolin Ke, Xiaodong Liu, Di He, Saurabh Tiwary, Tie-Yan Liu, Paul Bennett, Xia Song, and Jianfeng Gao. 2022. [METRO: efficient denoising pretraining of large scale autoencoding language models with model generated signals](#). *CoRR*, abs/2204.06644.
- Marco Baroni. 2022. [On the proper role of linguistically-oriented deep net analysis in linguistic theorizing](#). In Shalom Lappin, editor, *Algebraic systems and the representation of linguistic knowledge*, chapter 1, pages 5–22. Taylor and Francis, Abingdon-on-Thames.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, New York. Association for Computer Machinery – ACM.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrike, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. 2023. [Sparks of artificial general intelligence: Early experiments with gpt-4](#).
- Rosa Cao and Daniel Yamins. 2021. [Explanatory models in neuroscience: Part 1–taking mechanistic abstraction seriously](#). *arXiv preprint arXiv:2104.01490*.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [Electra: Pre-training text encoders as discriminators rather than generators](#). In *International Conference on Learning Representations*.
- The ManyBabies Consortium and Michael C. Frank et al. 2020. [Quantifying sources of variability in infancy research using the infant-directed-speech preference](#). *Advances in Methods and Practices in Psychological Science*, 3(1):24–52.
- Michael C Frank. 2023. [Large language models as models of human cognition](#).
- Richard Futrell and Roger P Levy. 2019. [Do RNNs learn human-like abstract word order preferences?](#) *Proceedings of the Society for Computation in Linguistics*, 2(1):50–59.
- Alison Gopnik, Andrew N Meltzoff, and Patricia K Kuhl. 1999. *The scientist in the crib: Minds, brains, and how children learn*. William Morrow & Co.
- Yuxian Gu, Zhengyan Zhang, Xiaozhi Wang, Zhiyuan Liu, and Maosong Sun. 2020. [Train no evil: Selective masking for task-guided pre-training](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6966–6974, Online. Association for Computational Linguistics.
- Curtis Hawthorne, Andriy Stasyuk, Adam Roberts, Ian Simon, Cheng-Zhi Anna Huang, Sander Dieleman, Erich Elsen, Jesse Engel, and Douglas Eck. 2019. [Enabling factorized piano music modeling and generation with the MAESTRO dataset](#). In *International Conference on Learning Representations*.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Richard Kunert, Raquel Fernández, and Willem Zuidema. 2011. [Adaptation in child directed speech: Evidence from corpora](#). In *Proceedings of the 15th SemDial Workshop on the Semantics and Pragmatics of Dialogue (Los Angeles)*, pages 112–119, Los Angeles, California, USA.
- Fred Lerdahl. 1996. [Calculating tonal tension](#). *Music Perception: An Interdisciplinary Journal*, 13(3):319–363.
- Tal Linzen and Marco Baroni. 2021. [Syntactic Structure from Deep Learning](#). *Annual Review of Linguistics*, 7(1):195–212.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. [Assessing the ability of LSTMs to learn syntax-sensitive dependencies](#). *Transactions of the Association for Computational Linguistics*, 4:521–535.
- Kyle Mahowald, Anna A. Ivanova, Idan A. Blank, Nancy Kanwisher, Joshua B. Tenenbaum, and Evelina Fedorenko. 2023. [Dissociating language and thought in large language models: a cognitive perspective](#).
- Gary F. Marcus. 1993. [Negative evidence in language acquisition](#). *Cognition*, 46:53–85.

- Aaron Mueller and Tal Linzen. 2023. [How to plant trees in language models: Data and architectural effects on the emergence of syntactic inductive biases](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11237–11252, Toronto, Canada. Association for Computational Linguistics.
- Howard Nicholas, Patsy M Lightbown, and Nina Spada. 2001. [Recasts as feedback to language learners](#). *Language learning*, 51(4):719–758.
- Isabel Papadimitriou and Dan Jurafsky. 2020. [Learning Music Helps You Read: Using transfer to study linguistic structure in language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6829–6839, Online. Association for Computational Linguistics.
- Steven T Piantadosi. 2023. [Modern language models refute chomsky’s approach to language](#). *Lingbuzz Preprint*, lingbuzz/007180.
- Ofir Press, Noah A. Smith, and Mike Lewis. 2021. [Shortformer: Better language modeling using shorter inputs](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5493–5505, Online. Association for Computational Linguistics.
- Nafis Sadeq, Canwen Xu, and Julian McAuley. 2022. [InforMask: Unsupervised informative masking for language model pretraining](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5866–5878, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Julian Salazar, Davis Liang, Toan Q. Nguyen, and Katrin Kirchhoff. 2020. [Masked language model scoring](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2699–2712, Online. Association for Computational Linguistics.
- Jessica F Schwab and Casey Lew-Williams. 2016. [Language learning, socioeconomic status, and child-directed speech](#). *Wiley Interdisciplinary Reviews: Cognitive Science*, 7(4):264–275.
- Michael Tomasello. 1992. [The social bases of language acquisition](#). *Social Development*, 1:67–87.
- Marten van Schijndel, Aaron Mueller, and Tal Linzen. 2019. [Quantity doesn’t buy quality syntax with neural language models](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5831–5837, Hong Kong, China. Association for Computational Linguistics.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems](#). In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*. Curran Associates Inc., Red Hook, NY, USA.
- Alex Warstadt and Samuel R. Bowman. 2022. [What artificial neural networks can tell us about human language acquisition](#).
- Alex Warstadt, Aaron Mueller, Leshem Choshen, Ethan Gotlieb Wilcox, Chengxu Zhuang, Juan Ciro, Rafael Mosquera, Adina Williams, Bhargavi Paranjabe, Tal Linzen, and Ryan Cotterell. 2023. [Findings of the 2023 BabyLM Challenge: Sample-efficient pretraining on developmentally plausible corpora](#). In *Proceedings of the 2023 BabyLM Challenge*. Association for Computational Linguistics (ACL).
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020a. [BLiMP: The benchmark of linguistic minimal pairs for English](#). *Transactions of the Association for Computational Linguistics*, 8:377–392.
- Alex Warstadt, Yian Zhang, Xiaocheng Li, Haokun Liu, and Samuel R. Bowman. 2020b. [Learning which features matter: RoBERTa acquires a preference for linguistic generalizations \(eventually\)](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 217–235, Online. Association for Computational Linguistics.
- Adriana Weisleder and Anne Fernald. 2013. [Talking to children matters: Early language experience strengthens processing and builds vocabulary](#). *Psychological science*, 24(11):2143–2152.
- Ethan Gotlieb Wilcox, Richard Futrell, and Roger Levy. 2022. [Using computational models to test syntactic learnability](#). *Linguistic Inquiry*, pages 1–88.
- Yian Zhang, Alex Warstadt, Xiaocheng Li, and Samuel R. Bowman. 2021. [When do you need billions of words of pretraining data?](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1112–1125, Online. Association for Computational Linguistics.

A Appendix

Table 3 shows the list of words selected for targeted MLM for each linguistic category, while age of acquisition results are presented in Table 4

| Category | Words |
|-------------------|--|
| S-V agreement | is, was, have, do, are, don't, were, has, does, isn't, doesn't, wasn't, haven't, aren't, weren't, hasn't |
| Quantifiers | all, some, more, any, little, many, much, most, every, both, each, few, enough, several, half, less, either, none, lots, neither, plenty |
| Filler gap | that |
| Modal verbs | can, would, will, could, should, may, must, might, shall |
| NPI licensing | not, only, also, really, probably, often, certainly, clearly |
| D-N agreement | this, these |
| Adverbs | never, always, maybe, probably, perhaps, certainly, absolutely, likely, possibly, definitely, surely, truly, constantly, forever, potentially, positively, undoubtedly, consistently, invariably, eternally, perpetually, dubiously, uncertainly |
| Anaphor agreement | himself, themselves, itself, herself |
| Animacy | people, man, men, family, person, father, mother, girl, woman, son, children, guy, friend, wife, boy, guys, human, member, friends, women, members, daughter, child, brother, boys, husband, girls, lady, parents, kids, king, sister, dad, mommy, daddy, player, students, doctor, president, captain, kid, mom, leader, officer, director, players, soldiers, teacher, god, student, sir, officers, judge, patient, brothers, families, mark, actor, ladies, singer, uncle, author, manager, gentleman, humans, lad, writer, sweetie, prince, lawyer, artist, mum, host, owner, guest, teachers, princess, scientists, guard, professor, artists, leaders, agent, assistant, patients, mama, workers, minister, boss, sons, criminal, partner, babies, citizens, adult, politician, gods, mayor, actress, principal, cousin, witness, driver, hero, governor, lord, doctors, authorities, maiden, suspect, victims, aunt, candidate, individuals, producer, champion, gentlemen, founder, enemies, sisters, winner, passenger, client, bride, priest, prisoners, pilot, inhabitants, ghost, chairman, nurse, guests, user, pirate, graduate, merchant, cats, victim, passengers, pirates, noble, agents, expert, parent, editor, grandma, officials, subjects, cops, maid, commander, policeman, writers, servants, academic, peasant, eldest, engineer, musician, devil, critics, users, creatures, twin, composer, personality, lads, followers, poet, adults, boyfriend, fellows, actors, ruler, judges, witch, daughters, lieutenant, musicians, servant, secretary, slave, priests, scholars, prisoner, visitors, residents, lover, cop, companion, knight, deputy, customers, tourist, guards, grandfather, journalist, architect, rival, kings, colleagues, farmers, owners, farmer,... |

Table 3: Words which were masked in targeted MLM in the 10M train set. For *Animacy* only words appearing over 100 times are shown in the table.

| Model | Overall | Nouns | Predicates | Function words |
|------------------|---------|-------|------------|----------------|
| RoBERTa-baseline | 2.06 | 1.99 | 1.85 | 2.65 |
| Lil-Bevo | 2.06 | 2.0 | 1.84 | 2.65 |
| Lil-Bevo-X | 2.05 | 1.99 | 1.85 | 2.59 |

Table 4: Age of Acquisition results