

Can training neural language models on a curriculum with developmentally plausible data improve alignment with human reading behavior?

Aryaman Chobey
Colgate University
achobey@colgate.edu

Oliver Smith
Colgate University
osmith@colgate.edu

Anzi Wang
Colgate University
awang1@colgate.edu

Grusha Prasad
Colgate University
gprasad@colgate.edu

Abstract

The use of neural language models to model human behavior has met with mixed success. While some work has found that the surprisal estimates from these models can be used to predict a wide range of human neural and behavioral responses, other work studying more complex syntactic phenomena has found that these surprisal estimates generate incorrect behavioral predictions. This paper explores the extent to which the misalignment between empirical and model-predicted behavior can be minimized by training models on more developmentally plausible data, such as in the BabyLM Challenge. We trained teacher language models on the BabyLM “strict-small” dataset and used sentence level surprisal estimates from these teacher models to create a curriculum. We found tentative evidence that our curriculum made it easier for models to acquire linguistic knowledge from the training data: on the subset of tasks in the BabyLM challenge suite evaluating models’ grammatical knowledge of English, models first trained on the BabyLM data curriculum and then on a few randomly ordered training epochs performed slightly better than models trained on randomly ordered epochs alone. This improved linguistic knowledge acquisition did not result in better alignment with human reading behavior, however: models trained on the BabyLM dataset (with or without a curriculum) generated predictions that were as misaligned with human behavior as models trained on larger less curated datasets. This suggests that training on developmentally plausible datasets alone is likely insufficient to generate language models capable of accurately predicting human language processing.

1 Introduction

The rapidly increasing success of neural language models has resulted in a corresponding increase the use of these models to model human neural and behavioral responses. This research direction has yielded mixed success — while the surprisal

estimates from these language models (i.e., the negative log probability of words given their preceding context) can certainly predict a wide range of neural and behavioral responses (Schrimpf et al., 2021), there are cases where surprisal estimates from these models generate quantitatively (Huang et al., 2023; Van Schijndel and Linzen, 2021) and even qualitatively (Arehalli and Linzen, 2020) incorrect predictions.

To what extent are these incorrect predictions a consequence of the fact that these models are trained on orders of magnitude more data than an average human is exposed to in their lifetime (Linzen, 2020)? Can training these models on more developmentally plausible datasets, such as in the BabyLM challenge (Warstadt et al., 2023), bridge the gap between empirical and predicted behavior? Does increased alignment with human behavior come at the cost of success on other NLP tasks? We explore these questions in this paper by training models on the the “strict-small” dataset of the BabyLM Challenge (~10M tokens) and evaluating the models on two types of tasks: first, tasks from the BabyLM challenge designed to test these models’ linguistic abilities; second, a large scale reading time dataset of syntactically complex sentences designed to evaluate models’ ability to capture aspects of human language processing (SAP benchmark; Huang et al., 2023).

Concretely, we explored whether training models on an easy-to-difficult curriculum (Elman, 1993) could result in improved performance on the BabyLM suite of challenge tasks and/or an improved fit to human reading behavior in the SAP Benchmark. To design the curriculum, we used the Cross-Review method (Xu et al., 2020): we trained teacher language models on different subsets of the training dataset and then generated sentence level surprisal estimates for held out sentences from each of the teacher models. For every sentence, the surprisal estimates from multiple teachers were av-

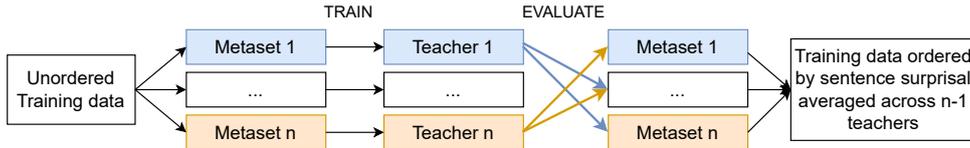


Figure 1: Schematic for how Cross-Review can be used to generate an easy-to-difficult order for the training dataset.

eraged together to compute a “difficulty” score, which was then used to generate an ordered sequence of training sentences (or “curriculum”).

To foreshadow our results, we found that models trained on our curriculum alone performed *worse* on the BabyLM suite of challenge tasks compared to models trained for the same number of steps without a curriculum. However, for a subset of the challenge tasks evaluating models’ grammatical knowledge, models which were trained on the curriculum followed by a few randomly ordered training epochs performed better than models trained on the randomly ordered epochs alone. This suggests that while training on the curriculum alone was not sufficient to acquire relevant linguistic knowledge, it might have induced useful biases in the models which made it easier for the models to acquire linguistic knowledge from the training data.

However, any useful biases that training on the the curriculum might have induced did not result in improved alignment with human reading behavior: models trained on the BabyLM data (with or without the curriculum) had nearly identical performance on the SAP benchmark to each other and to models trained on larger and less curated datasets. This result, along with prior work on training models on child directed speech (Yedetore et al., 2023), suggests that merely training on developmentally plausible data is likely insufficient for bridging the gap between human behavior and language-model predicted behavior.

2 Background

Curriculum learning (Bengio et al., 2009) refers to training models through a difficulty-based ordering of training examples (i.e. a curriculum), most often “starting small” (Elman, 1993) from easy examples before progressing to increasingly difficult sentences. In NLP, curriculum learning has been widely used for Machine Translation (e.g., Platanios et al., 2019), but has also been applied more recently to other natural language understanding tasks (Xu et al., 2020). For a survey see Soviany et al. (2022); Wang et al. (2021).

There are two steps involved in designing a curriculum: assigning a difficulty score to every training example (“difficulty measurer”) and using these difficulty scores to determine the order in which training examples are presented to the model (“training scheduler”) (Wang et al., 2021).

2.1 Difficulty measurer

Prior work exploring the efficacy of curriculum learning for NLU tasks has used a wide range of properties to compute sentence difficulty such as sentence length, word frequency (or rarity), tree depth, diversity and understandability (for a review, see Soviany et al., 2022). None of these properties by themselves can comprehensively capture what makes one sentence more difficult to process or acquire than another. For example, while long sentences are in general more difficult than short sentences, a shorter ambiguous sentence (“the horse raced past the barn fell”) is more difficult to process than a longer unambiguous one (“the horse which was raced past the barn is the same horse that fell”). Given the complex ways in which all of the individual properties can interact, a holistic way of combining these properties is likely necessary to generate good measures of sentence difficulty.

A natural way of combining these properties to compute a difficulty measure is to use a “teacher” language model to compute the predictability of words in a sentence: given some context, a good language model will assign lower probabilities to words that result in long continuations with infrequent words and structures and/or continuations that describe improbable or hard-to-understand events. Concretely, in this work we define difficulty of a sentence as the mean surprisal of words in the sentence, as given in equation 1, where \mathcal{D} is difficulty, L is the model being used to compute difficulty, s_k is the k -th sentence, and n is the number of words in s_k .

$$\mathcal{D}(s_k, L) = -\frac{1}{n} \sum_{i=0}^n \log P(w_i | w_0 \dots w_{i-1}, L) \quad (1)$$

There are two issues with estimating sentence difficulty in this manner. First, the difficulty estimates can be inaccurate if the teacher language model is trained on the same data for which difficulty scores are being computed. Second, the difficulty estimates can be affected by noisy idiosyncrasies if they are computed from just one teacher language model. To avoid these two issues, we use the Cross-Review method proposed by Xu et al. (2020). In this method each teacher is trained on a subset of the data, and then evaluated on all subsets other than the one it was trained on. Therefore if there are n teachers, there are $n-1$ difficulty scores for each sentence which can then be averaged together for a final difficulty score for the sentence (see Equation 2 and Figure 1).

$$\mathbf{D}(s_k) = \frac{1}{M} \sum_{m=1}^M \mathcal{D}(s_k, m) \quad (2)$$

2.2 Training scheduler

Given a training dataset E in which examples are ordered by difficulty and a training time step t , the training scheduler determines the subset of E that the model can be exposed to at t . At a broad level there are two types of schedulers: discrete and continuous (see Wang et al., 2021 for a more detailed taxonomy of training schedulers). In discrete schedulers, the training proceeds in stages with m training time steps; at all training time steps in a stage $t_i \dots t_{i+m}$, the model is exposed to the same subset of E .¹ In continuous schedulers on the other hand, the subset of E that the model is exposed to changes at every training time step.

In this work we use a continuous scheduler proposed by Platanios et al. (2019), in which the proportion of E that the model can be exposed to at t , $c_{root-p}(t)$, is given by the formula below, where T is the maximum number of training time steps and c_0 is the proportion of sentences that the model is exposed in the first time step:

$$c_{root-p}(t) = \min \left(1, \sqrt[p]{t \frac{1-c_0^p}{T} + c_0^p} \right) \quad (3)$$

Our primary reason for using the scheduler above is that it has only three hyperparameters: c_0 , T and p . The authors demonstrate that hyperparameters like warmup steps, which are normally

¹This is equivalent to saying that the at any given training stage, the model is trained on m epochs of a subset of E .

very highly tuned, do not have to be tuned with their scheduler. Given our compute limitations, hyperparameter tuning was infeasible, thus making this approach appealing.

3 Designing the curriculum

3.1 Datasets

We trained our random baselines and designed our curriculum using the datasets provided in the “strict-small” track of the BabyLM challenge. The data for this track was made up of 10 datasets, with a total of about $\sim 10M$ tokens and $\sim 920K$ sentences (where sentences were defined as sequences separated by a new line character). As specified in the BabyLM call for papers, the relative distribution of the ten datasets at the token level was intended to be developmentally plausible – for example, about 55% of all the tokens in the “strict-small” dataset comes from transcribed speech, and another 19% of the tokens come from stories (see Table 1).

While the BabyLM challenge datasets were constructed at the *token level*, we designed our curriculum at the *sentence level*, where we defined sentences as sequences separated by a new line character. We did this because it was more straightforward to sort the training dataset based on the difficulty of entire sentences; creating a token-level curriculum would require specifying an additional mechanism for ensuring that contextual integrity was maintained. The relative distribution of the ten datasets at the sentence level was very different from the relative distribution of tokens (see Table 1). Specifically, the proportion of more “complex” datasets (such as Wikipedia and Simple Wikipedia) was much lower at the sentence level than at the token level. We discuss the consequence of these distributional differences in § 7.

3.2 Computing sentence difficulty

As discussed in § 2.1, we used the Cross-Review method proposed by Xu et al. (2020) to compute the difficulty of every sentence in the training dataset. We divided the training dataset into five metaset, each with approximately the same number of tokens and number of sentences. Then, we used the neural-complexity codebase (van Schijndel and Linzen, 2018)² to train five LSTM teachers on each of these metaset.

Our LSTM teachers each had two hidden layers with 200 units in each layer. Training sentences

²<https://github.com/vansky/neural-complexity>

Dataset	Speech?	# tokens	Proportion	# sentences	Proportion
CHILDES	Yes	0.44 M	4%	80K	9%
BNC Spoken	Yes	0.84 M	9%	73.41 K	8%
Children’s book test	No	0.57 M	6%	26 K	3%
Children stories	No	0.34 M	3%	5.72 K	1%
Project Gutenberg	No	0.99 M	10%	91.81 K	10%
Open Subtitles	Yes	3.03 M	31%	470.89 K	51%
QED	Yes	1.03M	10%	91.91 K	10%
Simple Wikipedia	No	1.51 M	15%	48.80 K	5%
Switchboard	Yes	0.11 M	1%	11.09 K	1%
Wikipedia	No	0.99 M	10%	19.35 K	2%
Total		9.87 M		918.98 K	

Table 1: Number of tokens and sentences in each of the sub-datasets in the BabyLM “strict-small” datasets. The number of tokens are based on a BPE tokenizer we trained and we exclude tokens from lines with just a tab or space. Therefore the numbers are slightly different from those in the BabyLM call for papers.

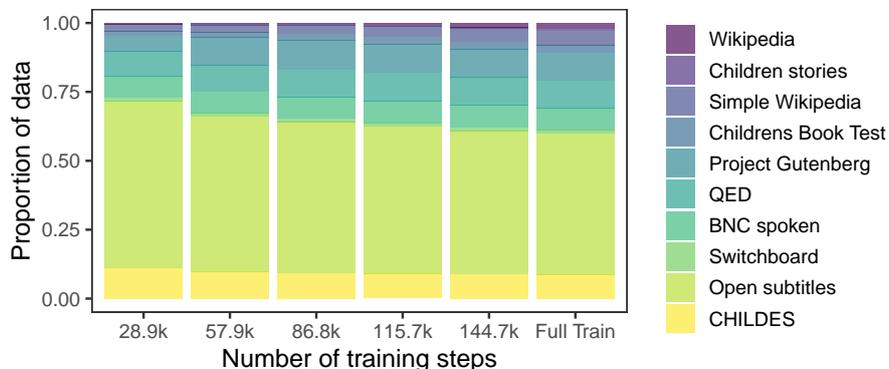


Figure 2: Proportion of sentences from each of the sub-datasets in our training curriculum for every 28937 training steps (i.e. equivalent to 1 epoch in the random models), as well as the proportions in the entire training dataset.

were pre-tokenized using a BPE tokenizer that we trained (described in § 4.2) and were passed to the teacher models in 20 batches. They were trained until their validation loss did not improve for three epochs, or until they reached 100 epochs. All teachers converged within 67 epochs, with the fastest teacher converging in 54 epochs.

We then evaluated each of the teacher LSTMs on all metaset except the one they were trained on, and then used the resulting surprisal values to compute the difficulty of every sentence in the training dataset (see Equation 2 and Figure 1).

Why use LSTM teachers? We trained LSTM language models instead of transformers because prior work has demonstrated that for datasets with 4 million tokens or less, such as our metaset, LSTM language models outperform their transformer counterparts (Hu et al., 2020), and therefore would make better “teachers”. Note, we did not use state-of-the-art language models as our teachers be-

cause of the constraints of the *strict-small track* of the BabyLM challenge.

3.3 Creating the training dataset

As discussed in § 2.2, we use the training scheduler proposed by Platanios et al. (2019) which has three hyperparameters (see Equation 3): the initial competence (c_0), the total number of training steps (T) and the root value (p). Following Platanios et al. (2019), we set the value of c_0 to 0.01. We set the value of T to be 150001 because our random baseline (described in § 4.3) achieved the highest validation perplexity after 144685 training steps (i.e., after 5 epochs).³ We set the value of p to be 10 after some experimentation because for values of p lower than that, the complex domains in our training dataset (such as Wikipedia) were very underrepresented (see Figure 5 in the Appendix). Then, for every batch, we sampled 32 sentences

³It was 150001 instead of 150000 because of an error.

from the subset of sentences that the model can be exposed to at the current time step as determined by Equation 3.

4 Training

4.1 Model architecture

For our target models we use the OPT 125M architecture (Zhang et al., 2022). This decoder-only transformer architecture consists of 12 layers with 12 attention heads each, an embedding size of 768 and a context length of 2048 tokens. We additionally use a final 0.2M layer with a causal language modeling head.

4.2 Tokenization and batching

Since the BabyLM challenge does not permit the use of pretrained tokenizer, we trained a BPE tokenizer on the training dataset with a vocabulary size of 50272 (the same as was used in the original OPT models). Like in the GPT-2 (Radford et al., 2019) tokenizer implementation, we do not significantly normalize or pre-tokenize the tokenizer training data. For the batching process, the tokenizer truncates sequences longer than 128 tokens, and returns the overflowing tokens as a separate sequence; only about 2% of our training examples were truncated. We used batch size of 32 with dynamic padding. The entire training dataset was divided into 28937 batches or training steps per epoch.

4.3 Model types

Random baseline: A randomly initialized OPT 125M model trained on our training dataset without any curriculum for up to 8 epochs. We present results from two baselines: the checkpoint after the 5th epoch (RandOPT 5ep; 144685 training steps) which had the best validation loss, and the last checkpoint (RandOPT 8 ep; 231496 training steps).

Curriculum only model: A randomly initialized OPT 125M model trained on our entire curriculum (CurrOPT; 150001 training steps).

Curriculum + Finetuning: The checkpoint of the CurrOPT model after it was trained on 144685 steps (i.e., same number of steps as the RandOPT 5ep model) further “finetuned” on the entire randomly ordered training dataset for upto 5 additional epochs. We present results from the checkpoint after 3 finetuning epochs (CurrOPT_ft 3ep; 231496 training steps, same as RandOPT 8ep) and the

checkpoint after 5 finetuning epochs (CurrOPT_ft 5ep; 289370 training steps).

4.4 Training procedure

We use an AdamW optimizer (Loshchilov and Hutter, 2017) with β_1 and β_2 set to 0.9 and 0.95 respectively. We use a weight decay and dropout of 0.1, and clip gradient norms at 1.0. For our random baseline we use a linear learning rate schedule and use a warmup of $\sim 5\%$ of our maximum training steps. As discussed in § 2.2 we do not use warmup for our curriculum models. Due to our considerably smaller pre-training corpus we do not implement the several mid-flight changes to learning rate and gradient clipping employed by Zhang et al. (as an adhoc response to training instability) over the course of their significantly longer training run.

5 Evaluation

We evaluate our models on the three challenge sets included in the BabyLM challenge – BLiMP (Warstadt et al., 2020a), (Super)GLUE (Wang et al., 2018, 2019) and MSGS (Warstadt et al., 2020b) — as well as on the SAP Benchmark (Huang et al., 2023).

BLiMP and BLiMP supplement The Benchmark of Linguistic Minimal Pairs (BLiMP) probes the linguistic knowledge that a language model encodes by measuring how often the model accurately assigns higher probabilities to words in minimally different grammatical and ungrammatical sentences. The original dataset contains minimal pairs for 12 different linguistic phenomena probing English morphology, syntax and semantics. The BabyLM challenge supplements this dataset with five additional linguistic phenomena targeting discourse level acceptability as well as other syntactic phenomena (such as question formation).

SuperGLUE The General language Understanding Evaluation (GLUE) benchmark and its successor SuperGLUE are challenge sets that are designed to evaluate models’ general purpose natural language understanding. The BabyLM challenge includes tasks from GLUE (COLA, SST2, MRPC, QQP, MNLI, QNLI, RTE), three tasks from SuperGLUE (BoolQ, RTE and WSC), as well as an additional task (Multimodal NLI). Unlike BLiMP which largely evaluates grammatical knowledge, the SuperGLUE tasks are designed to evaluate higher level linguistic abilities such as sentiment

analysis, inference, causal reasoning, coreference resolution, question answering, paraphrasing, etc.

MSGS The Mixed Signals Generalization Set is a diagnostic set used to evaluate how models solve an ambiguous classification task that can be solved using either linguistic features or surface features. The MSGS set contains five surface features and four linguistic features, resulting in 20 ambiguous classification tasks. There are also 9 control tasks to evaluate how well models can classify each of the features in an unambiguous context. The BabyLM challenge uses three linguistic features (syntactic position, syntactic construction, and syntactic category) and two surface features (lexical content and relative position), thus resulting in six ambiguous classification tasks.

SAP Benchmark The Syntactic Ambiguity Processing (SAP) benchmark is a large scaled reading time dataset for seven different types of syntactically complex sentences. Unlike the other datasets which measure models’ linguistic knowledge and ability, this dataset measures whether the models process information as humans do; specifically, whether models and humans are equally surprised by sentences that are grammatical but have complex and infrequent syntactic structures. The data processing pipeline of the SAP benchmark involves three steps: first, estimating *empirical* effects of interest using Bayesian mixed effects models; second, generating predicted reading times from language model surprisal (i.e. negative log probability) values and fitting mixed effects models to estimate *predicted* effects of interest;⁴ and third, comparing empirical and predicted effects of interest. The surprisal estimates and reading times are measured at specific *target* words and the following two spillover words. Further details about the different constructions are included in the Supplementary materials.

6 Results

6.1 What curriculum was learned?

The datasets with transcribed speech had the lowest average sentence difficulty scores. Even within transcribed speech, datasets with informal speech (such as child directed speech and subtitles) had lower average difficulty scores than datasets with

⁴SAP Benchmark uses Bayesian mixed effects models. We use linear mixed effects models because they are less resource intensive to fit and yield nearly identical model estimates.

more formal speech (such as BNC). Additionally, as expected the proportion of transcribed speech steadily decreased over time, as the proportion of written text increased. By the last “epoch”, the distribution of datasets was very similar to the true distribution (see Figure 2), suggesting that the cross-review method we used as our difficulty-measurer was effective, as was the root-10 training scheduler.

Agreement between LSTM teachers For any given sentence, there was a lot of variance in the surprisal estimates across the teachers: the average standard deviation was 113 bits of surprisal; the mean Spearman rank correlation between any two pairs of teachers was only 0.0009. This highlights the importance of averaging the surprisal estimates across different teachers to avoid over-fitting to idiosyncrasies of any particular teacher model.

Other difficulty measures Figure 7 plots the correlation between our difficulty measure computed using the cross-review method and two other simpler difficulty measures: average unigram frequency of the words in a sentence and sentence length. Our difficulty measure is moderately correlated with unigram frequency ($R = 0.27$, $p < 0.0001$) and highly correlated with sentence length ($R = 0.89$, $p < 0.0001$). We also predicted our difficulty measure as a function of unigram frequency and sentence length in a linear regression model and found that unigram frequency explains variance in our difficulty measure over and above sentence length, and together they explain most of the variance in the difficulty measure (adjusted R-squared = 0.93). This suggests that for the specific BabyLM datasets, using cross-review, while effective, might not be necessary: using faster-to-compute measures such as sentence length would have likely resulted in a comparable curriculum.

6.2 Training time

Since our difficulty measure was highly correlated with sentence length, in early stages of training the average sentence length in our curriculum was lower than the average sentence length in early epochs of model training without a curriculum. Since we dynamically padded our sequences, the model trained on our curriculum (CurrOPT) was initially trained on batches consisting of fewer total tokens than the model trained on the unordered data (RandOPT). As a result, in early stages of training, the time taken to train CurrOPT was less than half the amount of time taken to train RandOPT. As

Dataset	Task	RandOPT 5 ep	RandOPT 8 ep	CurrOPT	CurrOPT_ft 3 ep	CurrOPT_ft 5 ep
		114685 steps	231496 steps	150001 steps	231496 steps	289370 steps
BLiMP	Anaphor Agr.	75.72	86.71	70.35	72.59	75.05
	Agr. structure	66.09	67.85	67.8	70.44	70.04
	Binding	69.19	66.83	69.23	69.29	72.72
	Control/Raising	63.65	65.89	63.57	66.81	68.58
	D-N Agr.	72.33	74.64	72.05	76.58	78.6
	Ellipsis	52.71	52.89	53	61.66	55.43
	Filler-gap	72.84	73.5	72.74	74.4	75.18
	Irregular forms	82.39	71.04	82.34	80.97	81.27
	Island effects	52.06	57.21	57.1	64.13	62.48
	NPI licensing	47.46	41.59	38.76	45.52	48.25
	Quantifiers	55.69	64.4	52.81	67.03	67.34
	Subj-Verb Agr	63.92	64.43	58.84	64.73	65.55
BLiMP Supplement	Hypernym	50.58	49.19	48.95	47.91	47.21
	Congr. (easy)	48.44	51.56	50	53.12	53.12
	Congr. (tricky)	36.97	36.97	36.36	36.36	36.36
	Subj-Aux Inv.	84.92	86.53	72.55	84.58	85.02
	Turn taking	55	60.71	51.43	55.71	57.5
SuperGlue	COLA	3.2	9.35	3.2	9.77	8.91
	SST2	83.07	83.86	83.27	85.43	83.86
	MRPC	73.64	75.59	65.50	72.07	80.14
	QQP	76.86	77.27	74.78	77.33	76.83
	MNLI	65.75	67.07	64.63	65.18	65.3
	MNLI-MM	65.88	66.31	65.58	66.2	66.22
	QNLI	60.63	59.84	59.54	61.33	60.98
	RTE	51.51	45.46	47.48	53.54	48.49
	BoolQ	65.15	67.50	66.53	60.30	66.81
	MultiRC	55.53	48.85	56.74	46.55	47.54
WSC	56.63	61.45	61.45	61.45	61.45	
MSGs	MV lexical	-100	-100	-100	-100	-100
	MV position	-99.95	-98.39	-99.75	-88.76	-97.62
	SC lexical	0.18	-57.66	-58.62	-62.46	-69.88
	SC position	-62.68	-66.39	-62.82	-76.28	-62.88
	CR lexical	0	-4.17	-1.7	-2.4	-1.2
	CR position	-69.49	-95.59	-87.47	-70.38	-98.53

Table 2: Results for the tasks included in the BabyLM challenge set. Most of the numbers in the table indicate accuracy except for the following cases: MSGS tasks and COLA numbers are Matthew’s Correlation Coefficient; MRPC and QQP numbers are F1 scores. Light green cells indicate cases in which CurrOPT and CurrOPT_ft 3ep performed better than their random counterpart trained on the same number of steps: RandOPT 5ep and RandOPT 8ep respectively. Similarly, red cells indicate cases in which CurrOPT and CurrOPT_ft 3ep perform worse. Orange cells indicate tasks in which one of the random models ultimately had the best performance. Teal cells indicate tasks in which training the random model on more epochs led to *worse* performance. Bolded numbers indicate the best performance in a task across all five models. For MSGS we interpret “best performance” as having the weakest surface bias (i.e., the least negative numbers).

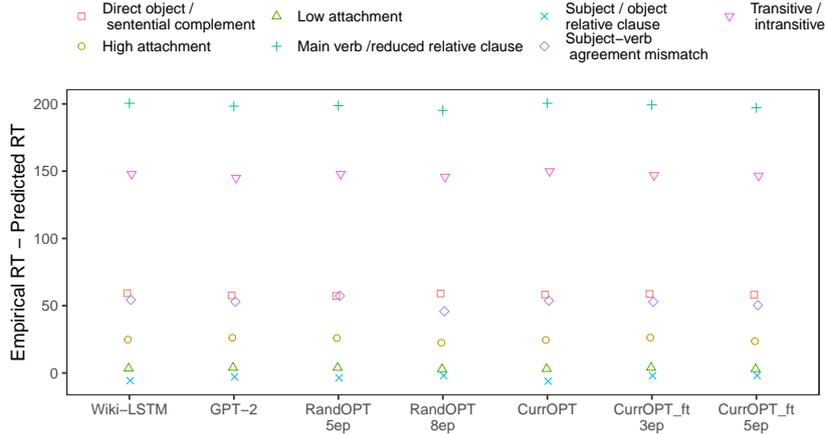


Figure 3: Difference between empirical and model-predicted reading times for the different constructions in the SAP Benchmark. Wiki-LSTM and GPT-2 estimates were from the original paper.

the curriculum progressed, the number of tokens in each batch of CurrOPT approached those in RandOPT causing the training time for CurrOPT to be similar to that of RandOPT.

6.3 BabyLM challenge tasks

On almost all of the tasks, the performance of the model trained on our curriculum (CurrOPT) was worse than the random baseline trained on fewer training examples (RandOPT 5ep). However, when we continued to train CurrOPT on more epochs of the entire training data, the resulting model (CurrOPT_ft 3ep) performed better than the random baseline trained on the same number of training examples (RandOPT 8ep) on some tasks (see Table 2). Specifically, CurrOPT_ft 3ep performed better than RandOPT 8ep on most tasks that evaluated models’ knowledge of English grammar (e.g., BLiMP, COLA). However, additional training did not seem to help the curriculum models’ performance on tasks that required specific lexical knowledge (e.g., irregular forms and hypernyms) or on tasks that required the model to learn more factual information (e.g., MNLI, MNLI-MM and BoolQ). Taken together these results suggest that while training on our curriculum by itself is insufficient to impart the necessary grammatical knowledge, it might induce biases in the model that make it easier for the model to acquire this knowledge from training data. However, there may be limits to the usefulness of these induced biases: training on our curriculum seemed to have some negative impact on the models’ ability to acquire nuanced lexical or factual information required to solve more complex tasks

like inference or question answering.

6.4 SAP Benchmark

We compared reading times predicted from the surprisal estimates of each of our models, as well as two baselines that were used in the original paper (GPT-2 (Radford et al., 2019) and an LSTM model trained on Wikipedia) to the empirical reading times. The difference between predicted and empirical reading times is nearly identical across all models, and very high (greater than 25 ms) for five out of the seven constructions (see Figure 3). This difference is not just a result of an incorrect conversion from surprisal to RTs — we observe qualitatively similar patterns when we look at raw surprisal values (see Figure 8 in the Appendix). Thus training on developmentally plausible data (with or without a curriculum) does not result in more human-like processing compared to models trained on less curated written text from the internet. This result aligns with the finding that training on child directed speech does not result in human-like generalization (Yedetore et al., 2023). Taken together these results suggest that merely modifying the training data of language models is unlikely to result in better cognitive models of human language acquisition and processing.

7 Discussion

In this paper we explored whether training on a developmentally plausible dataset can improve alignment with human behavior, and whether the improved alignment (if any) comes at the cost of performance on other NLP tasks evaluating different

aspects of linguistic competence. We trained models with and without a curriculum on the BabyLM “strict-small” dataset and evaluated them on the BabyLM suite of evaluation tasks as well as on a large scale benchmark of reading behavior for syntactically complex sentences (SAP benchmark).

Drawing on prior work on curriculum learning, we created an easy-to-difficult ordering of the sentences in the training dataset using surprisal values from LSTM teacher language models in a Cross-Review paradigm (Xu et al., 2020), and then used this ordering with a root-10 scheduler (Platanios et al., 2019) to design the training curriculum. This learned curriculum aligned with intuitive expectations for our curriculum — for example, the proportion of transcribed speech decreased over time, whereas the proportion of written text increased.

An OPT125M causal language model trained on our curriculum (CurrOPT) performed *worse* on most of the tasks in the BabyLM challenge set compared to baselines trained without a curriculum, suggesting that the models were unable to acquire relevant linguistic knowledge from the curriculum alone. Continuing to train CurrOPT on epochs of randomly ordered training data improved performance on most tasks targeting grammatical knowledge, but not on tasks that required more fine-grained knowledge about lexical or factual content.

Why did training on the curriculum lead to worse performance on some tasks? Domains with complex sentences (e.g., Wikipedia) were underrepresented in our curriculum because of our sentence level curriculum: domains like Wikipedia had fewer but longer sentences, and were therefore less likely to be sampled than sentences from domains with many short sentences (e.g., Open Subtitles). As a consequence there might not have been enough signal in the training data for the models to acquire factual information (which might explain their poor performance on tasks like MNLI and BoolQ) or nuanced lexical representations (which might explain their poor performance on tasks like irregular forms and hypernyms).

Can training on developmentally plausible data improve alignment with human behavior? Crucial to our question, we found that our models which were trained on developmentally plausible data (with or without a curriculum) had nearly identical performance to models trained on less curated larger datasets — all of the models severely under-

predicted the magnitude of processing difficulty in syntactically complex sentences. This suggests that training on developmentally plausible data alone is likely insufficient to bridge the gap between human and model-predicted behavior.

Limitations and future work All of the performance increases that we’ve discussed were very modest and based on just one model architecture. Therefore further work with additional random runs of the model is required to ensure that the improvements in performance were not just random noise. Similarly repeating the experiments with different architectures for the target and teacher models can shed light on the generalizability of our conclusions. In a similar vein, the conclusions about SAP benchmark results also need to be validated in future work. Specifically, it is necessary to more carefully define what “developmentally plausible” means, develop concrete hypotheses about why training on specific datasets might result in better alignment with reading behavior, and test these hypotheses with controlled experiments.

8 Conclusion

We designed a surprisal-based curriculum using the developmentally plausible data in the BabyLM strict-small dataset. We found that a model which was first trained on this curriculum and then trained on several additional epochs of the unordered training dataset performed slightly better than a random baseline trained on the same number of examples across a range of NLP tasks. When these models were evaluated on the SAP benchmark, their performance was nearly identical to each other and to that of models trained on larger and less curated datasets. This suggests that merely altering the training data to be more developmentally plausible is unlikely to improve alignment with human behavior.

Acknowledgements

This work was supported in part through the Colgate University’s ITS HPC resources, services, and staff expertise. We would also like to thank Tom McCoy and Suhas Arehalli and NYU’s ITS HPC resources for their assistance with evaluating our models on the BabyLM evaluation tasks, as well as Suhas Arehalli and Forrest Davis for their valuable feedback.

References

- Suhas Arehalli and Tal Linzen. 2020. Neural language models capture some, but not all, agreement attraction effects. In *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*, pages 370–376.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48.
- Jeffrey L Elman. 1993. Learning and development in neural networks: The importance of starting small. *Cognition*, 48(1):71–99.
- Jennifer Hu, Jon Gauthier, Peng Qian, Ethan Wilcox, and Roger Levy. 2020. A systematic assessment of syntactic generalization in neural language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1725–1744, Online. Association for Computational Linguistics.
- Kuan-Jung Huang, Suhas Arehalli, Mari Kugemoto, Christian Muxica, Grusha Prasad, Brian Dillon, and Tal Linzen. 2023. Surprisal does not explain syntactic disambiguation difficulty: evidence from a large-scale benchmark.
- Tal Linzen. 2020. How can we accelerate progress towards human-like linguistic generalization? In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Seattle, Washington. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Emmanouil Antonios Platanios, Otilia Stretcu, Graham Neubig, Barnabas Poczos, and Tom M Mitchell. 2019. Competence-based curriculum learning for neural machine translation. *arXiv preprint arXiv:1903.09848*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Martin Schrimpf, Idan Asher Blank, Greta Tuckute, Carina Kauf, Eghbal A Hosseini, Nancy Kanwisher, Joshua B Tenenbaum, and Evelina Fedorenko. 2021. The neural architecture of language: Integrative modeling converges on predictive processing. *Proceedings of the National Academy of Sciences*, 118(45):e2105646118.
- Petru Soviany, Radu Tudor Ionescu, Paolo Rota, and Nicu Sebe. 2022. Curriculum learning: A survey. *International Journal of Computer Vision*, 130(6):1526–1565.
- Marten van Schijndel and Tal Linzen. 2018. A neural model of adaptation in reading. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4704–4710, Brussels, Belgium. Association for Computational Linguistics.
- Marten Van Schijndel and Tal Linzen. 2021. Single-stage prediction models do not explain the magnitude of syntactic disambiguation difficulty. *Cognitive science*, 45(6):e12988.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. Superglue: A stickier benchmark for general-purpose language understanding systems. *Advances in neural information processing systems*, 32.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Xin Wang, Yudong Chen, and Wenwu Zhu. 2021. A survey on curriculum learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):4555–4576.
- 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.
- Benfeng Xu, Licheng Zhang, Zhendong Mao, Quan Wang, Hongtao Xie, and Yongdong Zhang. 2020. Curriculum learning for natural language understanding. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6095–6104.

Aditya Yedetore, Tal Linzen, Robert Frank, and R. Thomas McCoy. 2023. [How poor is the stimulus? evaluating hierarchical generalization in neural networks trained on child-directed speech](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9370–9393, Toronto, Canada. Association for Computational Linguistics.

Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. [Opt: Open pre-trained transformer language models](#). *arXiv preprint arXiv:2205.01068*.

Main verb / reduced relative clause garden path

The little girl fed the lamb **remained** relatively calm ...

The little girl who was fed the lamb **remained** relatively calm ...

Direct object / sentential complement garden path

The little girl found the lamb **remained** relatively calm ...

The little girl found that the lamb **remained** relatively calm ...

Transitive / intransitive garden path

When the little girl attacked the lamb **remained** relatively calm ...

When the little girl attacked, the lamb **remained** relatively calm ...

Subject / object relative clause

The bus driver that **the** kids followed ...

The bus driver that followed **the** kids ...

Relative clause modifiers recent noun (low attachment)

Janet charmed the executives of the assistant who **decides** almost everything ...

Janet charmed the executive of the assistant who **decides** almost everything ...

Relative clause modifiers distant noun (high attachment)

Janet charmed the executive of the assistants who **decides** almost everything ...

Janet charmed the executive of the assistant who **decides** almost everything ...

Subject-verb agreement mismatch

Whenever the nurse calls, the doctors **stops** working immediately ...

Whenever the nurse calls, the doctor **stops** working immediately ...

Figure 4: Figure and caption adapted from (Huang et al., 2023). Each sentence pair illustrates a construction tested in SAP Benchmark. An effect of interest is defined as the difference in reading times associated with a disambiguating or ungrammatical word, marked in green, minus the reading time associated with that same word in a context where it is grammatical and does not disambiguate the structure of the sentence, marked in turquoise.

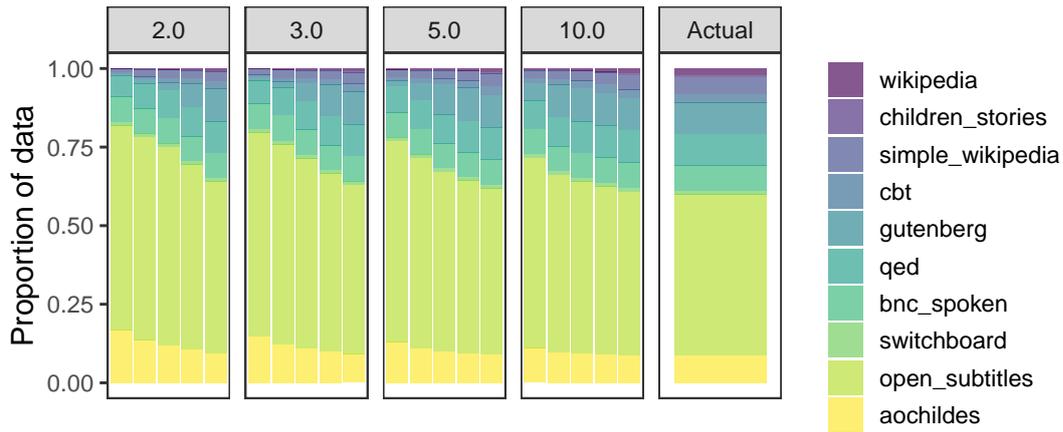


Figure 5: Proportion of sentences from each of the sub-datasets in training curricula with different root values for every 28937 training steps (i.e. equivalent to 1 epoch in the random models), as well as the proportions in the entire training dataset. The curriculum used this is paper is root 10.

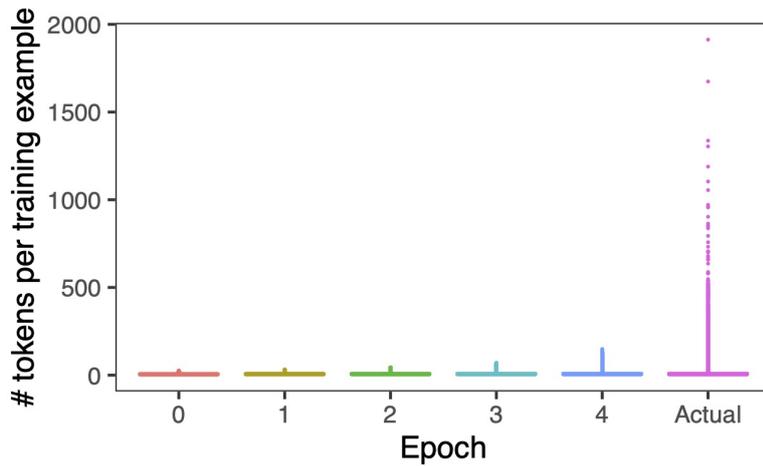


Figure 6: Caption

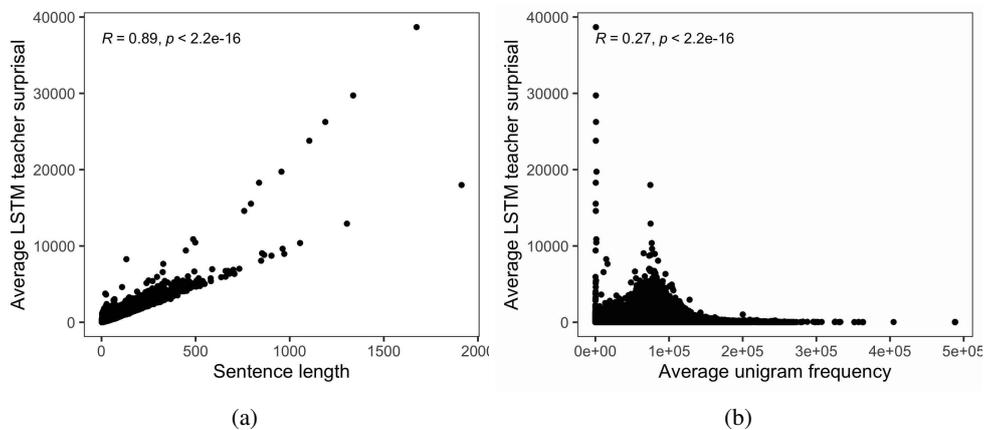


Figure 7: Relationship between the order average LSTM teacher surprisal and other difficulty measures. R values indicate the Spearman rank correlation coefficients.

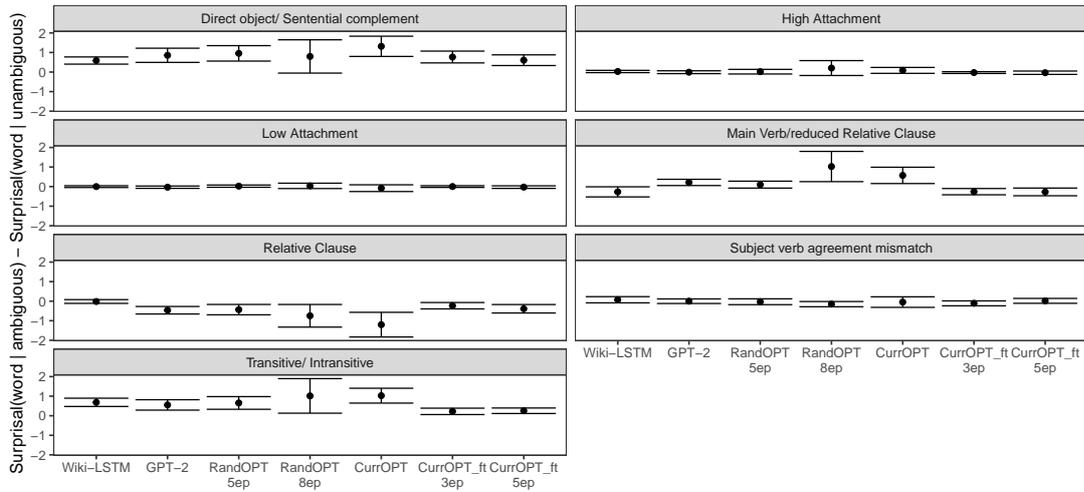


Figure 8: Difference in surprisal value at the target word in ambiguous and unambiguous sentences averaged across all sentences in a construction. Error bars represent 2 standard errors from the mean.

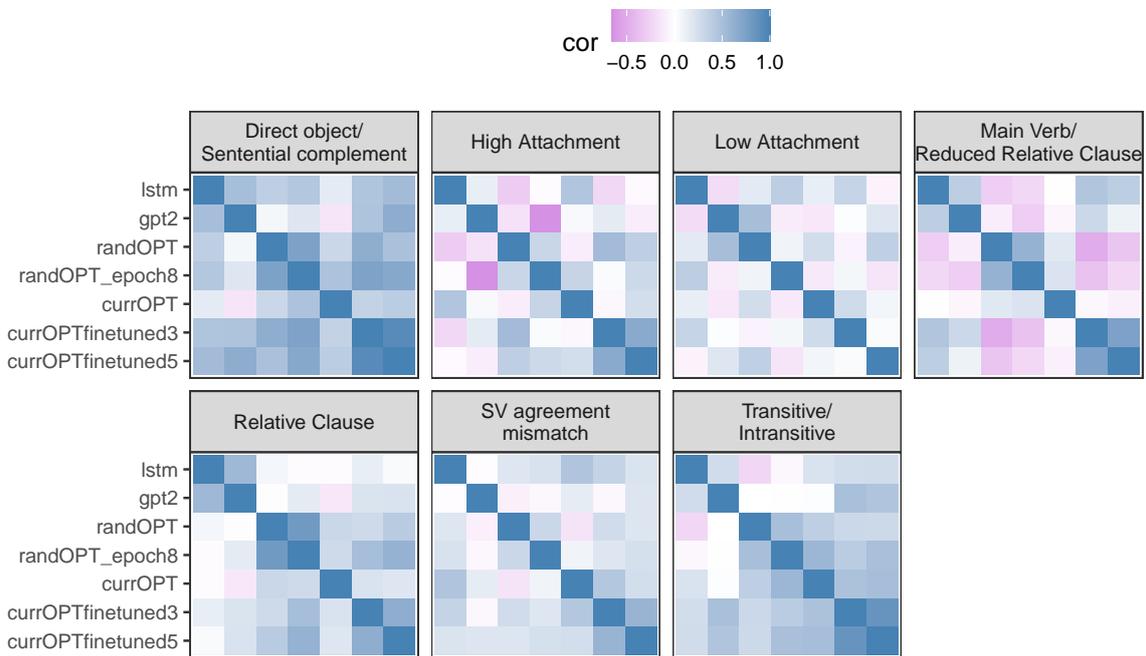


Figure 9: Correlation between model surprisal estimates at the item level. For each model, an item level difference in surprisal at the target word in the ambiguous and unambiguous conditions was computed. This item-level difference was correlated between models using Pearson's correlation. The relatively weak correlations suggest that even though the models have the same aggregate behavior on the SAP benchmark, their behavior differs at the item level.