

Mean BERTs make erratic language teachers: the effectiveness of latent bootstrapping in low-resource settings

David Samuel

University of Oslo, Language Technology Group

Abstract

This paper explores the use of latent bootstrapping, an alternative self-supervision technique, for pretraining language models. Unlike the typical practice of using self-supervision on discrete subwords, latent bootstrapping leverages contextualized embeddings for a richer supervision signal. We conduct experiments to assess how effective this approach is for acquiring linguistic knowledge from limited resources. Specifically, our experiments are based on the BabyLM shared task, which includes pretraining on two small curated corpora and an evaluation on four linguistic benchmarks.

1 Introduction

All modern language models are trained with a general self-supervised learning (SSL) paradigm (Radford et al., 2018; Devlin et al., 2019; Raffel et al., 2020). Recently, the field of visual representation learning has seen a growing usage of self-supervision on *latent embeddings* (Grill et al., 2020; Chen et al., 2020; Chen and He, 2020; Assran et al., 2023). While this type of self-supervision has been recently proposed as an integral part of a human-like machine intelligence system (LeCun, 2022), language models are still mostly self-supervised on hard targets, typically on subword tokens.

The concept of *latent bootstrapping* (Grill et al., 2020) offers a promising alternative, as the latent vectors provide a deep and semantically rich representation of the input. This, in turn, delivers a more valuable supervision signal compared to the conventional method of supervision on discrete subword indices. Data2vec (Baevski et al., 2022) showed that latent bootstrapping performs on par with traditional self-supervised language modeling when pretrained on a large text corpus. We argue that, intuitively, the rich training signal from contextualized embeddings should be particularly effective in low-resource data settings.

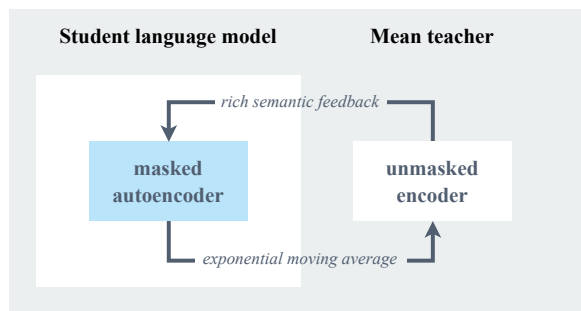


Figure 1: The self-supervision feedback loop of latent bootstrapping: a student model improves by aligning with its teacher’s latent outputs and the teacher improves by maintaining the exponential moving average of the student.

In this paper, our aim is to test this hypothesis and identify possible drawbacks of the bootstrapping method. We base our experiments on the *BabyLM challenge* (Warstadt et al., 2023b), a shared task that uses two carefully curated, sample-efficient pretraining corpora, mimicking the English language exposure to young children. In addition, this challenge employs four benchmarks to evaluate different aspects of linguistic knowledge and understanding learned by language models.

We introduce BootBERT, a novel masked autoencoder language model (Meng et al., 2023) that harnesses latent bootstrapping (Grill et al., 2020) between a mean teacher (Tarvainen and Valpola, 2017) and its student. Through a positive feedback loop, the student and the teacher iteratively learn from each other, as illustrated in Figure 1. The student is trained to match the teacher’s outputs while the *mean teacher* is defined as the exponential moving average of the student. Once pretraining is complete, only the student language model is used for evaluation and the teacher is discarded. We assess its performance on the BabyLM challenge, contrasting it with conventional language models. The source code and pretrained models are available online at <https://github.com/lgtoslo/boot-bert>.

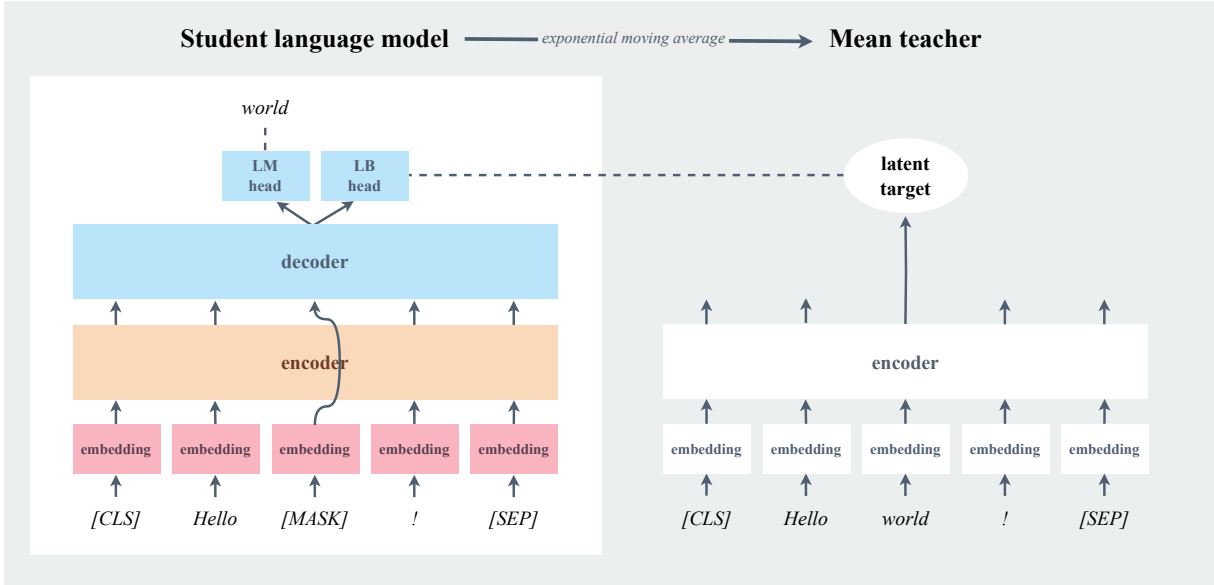


Figure 2: A detailed overview of the self-supervised feedback loop. The left side illustrates the student language model, a masked autoencoder network, that targets two training objectives: 1) conventional masked language modeling, aiming to predict the masked token (e.g., the word ‘world’), and 2) aligning the contextualized embedding of the masked tokens to their unmasked counterparts. The embeddings for the unmasked tokens are produced by a mean teacher network (on the right), computed as an exponential moving average of the student parameters.

2 Method

In this section, we outline our proposed model, BootBERT, delving into its neural architecture and the latent bootstrapping training objective. In order to allow for language-modeling-based evaluation, the bootstrapping objective operates alongside conventional masked language modeling. The diagram in Figure 2 illustrates the general idea of this approach.

Masked autoencoder architecture. BootBERT diverges slightly from the standard ‘encoder-only’ architecture often found in masked language models (Devlin et al., 2019). Instead, following the method of Meng et al. (2023), we employ a masked autoencoder (MAE; He et al., 2022) framework for the text domain. This approach distinguishes the *encoding* of contextualized embeddings from the *decoding* of masked subwords. These two functionalities are separated by dividing the model into an encoder and a decoder module, as illustrated in Figure 2 on the left.

The encoder’s role is to create a bidirectional contextualized embedding of input tokens. Unlike traditional masked language models, the encoder does not process any [MASK] tokens, thus eliminating the need to allocate parameters for representing them (Meng et al., 2023).

The [MASK] tokens are processed and denoised

by the decoder module. The decoder is supplied with the full input – the unmasked tokens are represented by their contextualized embeddings (provided by the encoder) and the masked tokens are represented by a static [MASK] embedding. Note that the decoder in this type of model is bidirectional and purely self-attentive, differing from the original definition of a transformer decoder by Vaswani et al. (2017).

Teacher-student feedback loop. Conceptually, the training process can be divided into optimization of a student model and optimization of a teacher model. Here, the masked student autoencoder model is trained to match the contextualized embeddings of the *unmasked* tokens, produced by the mean teacher network. In line with Tarvainen and Valpola (2017), the teacher parameters ϕ are not optimized via gradient descent, but rather through a slow exponential moving average (EMA) of the student parameters θ :

$$\phi = \tau\phi + (1 - \tau)\theta.$$

This moving average not only stabilizes the latent targets but also prevents representation collapse (Grill et al., 2020).

Loss. We optimize two objectives during training the student model: a traditional masked language

modeling objective with hard targets, symbolized by \mathcal{L}_{LM} , and a latent bootstrapping objective using teacher’s latent targets \mathcal{L}_{LB} . The final loss function combines these objectives with a weighted sum:

$$\mathcal{L} = \mathcal{L}_{LB} + \beta\mathcal{L}_{LM}.$$

Here, \mathcal{L}_{LM} is calculated simply as negative log-likelihood of the true targets. Its purpose is two-fold: allowing for a MLM-based evaluation (for example BLiMP), and preventing representation collapse of unconstrained latent bootstrapping (Grill et al., 2020).

The second objective is computed as a smooth L1 loss between student predictions y_s and teacher’s contextualized embeddings y_t . This works mostly like a standard mean-squared error but prevents exploding gradients from outliers (Girshick, 2015):

$$\mathcal{L}_{LB}(y_t, y_s) = \begin{cases} 0.5(y_t - y_s)^2 & |y_t - y_s| \leq 1 \\ |y_t - y_s| - 0.5 & \text{otherwise.} \end{cases}$$

LTG-BERT transformer backbone. As for more low-level architectural and training choices, we adopt the approach of LTG-BERT by Samuel et al. (2023a). This method was optimized for low-resource masked language modeling on a similar corpus to the corpora provided in BabyLM. The key improvements of the LTG-BERT transformer architecture include the use of the NormFormer layer normalization (Shleifer and Ott, 2022), an alternative disentangled attention mechanism with relative positions (He et al., 2021) and gated-linear activation function (GEGLU; Shazeer, 2020); as illustrated in Figure 3. On top of these architectural changes, the authors also employ masking of random subword spans (Joshi et al., 2020). More details about these choices can be found in Samuel et al. (2023a).

3 Experiments

The main goal of this paper is to evaluate how well language models trained with latent bootstrapping acquire language and if it makes a viable training objective for language representation learning. We base the experiments on the BabyLM challenge (Warstadt et al., 2023b). First, we describe the pretraining process of two BabyLM tracks and second, the evaluation of pretrained models using the BabyLM evaluation pipeline.

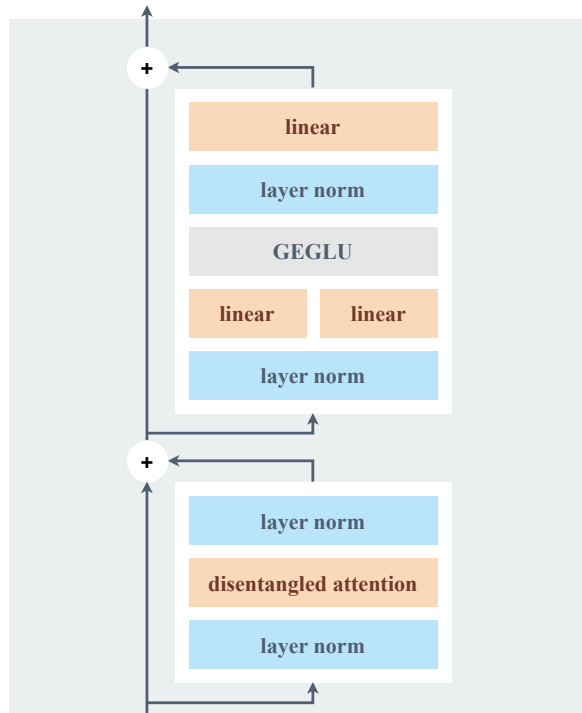


Figure 3: We base our model on LTG-BERT. This simplified diagram shows one layer from that transformer architecture, it illustrates the self-attention module (bottom) and the feed-forward module (top). Both modules utilize a modified NormFormer-like layer normalization placement and the feed-forward module contains a gated-linear activation function.

BabyLM challenge. This challenge provides a share ground for experiments on small-scale language modeling. It consists of three tracks: STRICT, STRICT-SMALL and LOOSE. For the first two tracks, the submissions have to be pretrained solely on a fixed corpus provided by the organizers. This corpus contains about 100M words in the STRICT track and about 10M words in the STRICT-SMALL track. As for the LOOSE track, the submissions are still limited to pretrained on 100M words, but this data can come from any source and the models can utilize an unlimited amount of non-linguistic data in addition. As detailed in Section 3.2, the submissions are compared on a shared evaluation set consisting of syntactic and natural language understanding tasks.

3.1 Pretraining

The pretraining is done on corpora provided by the BabyLM challenge. These texts are curated specifically to be of the same type and quantity that children learn from. Thus, it allows us to assess (to some degree) whether latent bootstrapping is a

more plausible cognitive model of human language acquisition (Warstadt et al., 2023b).

Training corpus. Specifically, we consider the STRICT and STRICT-SMALL tracks and pretrain the models on their respective 100-million-word and 10-million-word corpora. Both datasets contain child-directed speech, transcribed speech, children’s books and Wikipedia, among other sources. The content of these datasets is detailed in Appendix B, together with our simple preprocessing pipeline, which unifies the typographical features of the BabyLM subcorpora.

Pretraining process. Generally speaking, we adopt the training recipe of LTG-BERT (Samuel et al., 2023a), which was optimized for pretraining on another low-resource 100 million English corpus. The pretraining process is the same for both tracks, except for using a smaller vocabulary and a smaller model for the STRICT-SMALL track.

As for the STRICT track, we use a BASE-size language model – 12 encoder layers and 4 decoder layers with hidden size of 768 and with 12 attention heads. We train a case-sensitive WordPiece tokenizer (Wu et al., 2016) with a vocabulary size of $2^{14} = 16\,384$, using solely texts from the STRICT corpus. As per Samuel et al. (2023a), we pretrain the models with $\frac{1}{2}$ of the BERT training budget, as it has been shown to be sufficient for a relatively small 100-million-word corpus. The tokens are masked with continuous span masking (Joshi et al., 2020; Raffel et al., 2020). In particular, the masks are iteratively sampled until 15% of tokens are masked and the length of each span is sampled from the geometric distribution $\text{Geo}(p)$, with $p = \frac{1}{3}$.

The STRICT-SMALL track is tackled by a SMALL-size language model – 12 encoder layers and 4 decoder layers with hidden size of 384 and with 6 attention heads. The subword vocabulary is reduced to $2^{12} = 4\,096$ items.¹

The full list of hyperparameters and implementation details are provided in Appendix C and in the released source code.²

¹This choice is selected according to Gowda and May (2020) who recommend to ‘...use the largest possible vocabulary such that at least 95% of classes have 100 or more examples in training.’

²<https://github.com/ltgoslo/boot-bert>

3.2 Evaluation

We utilize the language modeling benchmark suite from the BabyLM challenge (Gao et al., 2021; Warstadt et al., 2023b),³ which relies on three conceptually different evaluation tasks:

1. The GLUE and SuperGLUE datasets test the ability of a pretrained model to adapt to various language understanding tasks.
2. BLiMP and BLiMP supplement tasks test the affinity of a model towards grammatical sentences in a completely zero-shot manner.
3. MSGS measures how much does a pretrained model prefer linguistic generalizations (over surface ones) during finetuning.

We further elaborate on each of these evaluation suites below.

(Super)GLUE benchmark. General Language Understanding Evaluation benchmarks (GLUE and SuperGLUE; Wang et al., 2018, 2019) are arguably the most common ways of evaluating the language-understanding and transfer-learning capabilities of language models. The BabyLM challenge uses a subset of 10 (Super)GLUE tasks, detailed in Appendix F. We employ the standard way of finetuning masked language models on these datasets, as introduced in BERT (Devlin et al., 2019). More details about the finetuning processes are given in Appendix C.

As we use the BabyLM version of GLUE, our results cannot be directly compared with previous literature – the dataset samples are filtered to not contain out-of-vocabulary words and some of the employed metrics differ from the original recommendations (Wang et al., 2018, 2019). We opted to adhere to the BabyLM version to be compatible with other works in this challenge. However, in order to reliably compare our models, we decided to depart from BabyLM and to divide the training set in 90:10 ratio into a new training and development split; the former validation set is then used as a held-out split.⁴

BLiMP. When using any finetuning approach, it is unclear how to disentangle the innate language

³<https://github.com/babylm/evaluation-pipeline>

⁴The BabyLM pipeline unfortunately uses identical validation and test sets, which might yield overly optimistic results due to overfitting during hyperparameter optimization.

understanding from the knowledge learned during the second-stage supervised finetuning (Belinkov, 2022). In contrast, the Benchmark of Linguistic Minimal Pairs (BLiMP; Warstadt et al., 2020a) attempts to measure the linguistic knowledge of a language model in a zero-shot manner – without any additional training. Each pair of sentences in BLiMP differs minimally on the surface level, but only one of the sentences is grammatically valid. We can use the intrinsic ability of language models to assign a probability to every sentence and test how often a language model assigns a higher probability to the correct sentence (Wang and Cho, 2019; Salazar et al., 2020).

As detailed in Appendix A, the results on BLiMP greatly depend on temperature scaling (Guo et al., 2017a). Thus, to fairly compare different types of language models, we employ an alternative approach to evaluating BLiMP: we report the accuracies that are achieved with the optimal temperature for every language model; the reasoning is explained in Appendix A.

The BabyLM challenge also comes with an additional ‘BLiMP supplement’ held-out set with five additional diagnostic tasks. To comply with the held-out spirit of these tasks, we keep the temperature values calibrated for BLiMP, even though this results in suboptimal performance (Appendix A).

MSGS. The diagnostic dataset called Mixed Signals Generalization Set (MSGS; Warstadt et al., 2020b) measures whether a pretrained model prefers linguistic or surface generalizations. The experiments follow *the poverty of the stimulus design* (Wilson, 2006) – to first finetune a model on ambiguous data (consistent with both linguistic and surface explanations) and then test it on non-ambiguous data to see if it prefers the linguistic generalization.

We use the filtered MSGS datasets with no *in-oculation* in the training set, as provided by the BabyLM challenge. Similarly to (Super)GLUE, we avoid the BabyLM approach that validates and tests on the same split – instead, to obtain a reliable comparison, we roughly follow the original work (Warstadt et al., 2020b) and use three learning rates: ($1 \cdot 10^{-5}$, $2 \cdot 10^{-5}$, and $3 \cdot 10^{-5}$), five random seeds, batch size of 16 and finetune for 5 epochs without early-stopping; then we report the mean and standard deviation statistics on the 6 non-ambiguous and non-control test datasets, measuring the Matthew’s correlation coefficient

Model	GLUE	MSGS	BLiMP	Supplement
STRICT (100M words)				
OPT _{125m}	73.0 \pm 3.9	-44.4 \pm 8.5	77.8	67.5
RoBERTa _{base}	74.3 \pm 0.6	-66.4 \pm 26.6	76.2	63.8
T5 _{base}	75.3 \pm 1.1	-56.5 \pm 6.7	83.6	71.8
LTG-BERT _{base}	77.8 \pm 1.4	-43.2 \pm 11.0	87.2	77.6
BootBERT _{base}	79.2 \pm 1.5	-67.9 \pm 12.6	86.3	72.2
STRICT-SMALL (10M words)				
OPT _{125m}	68.3 \pm 3.3	-63.8 \pm 9.6	69.2	60.2
RoBERTa _{base}	72.2 \pm 1.9	-66.7 \pm 11.9	68.1	60.5
T5-base	64.7 \pm 1.3	-68.4 \pm 7.1	59.9	48.6
LTG-BERT _{small}	74.5 \pm 1.5	-42.6 \pm 34.8	80.9	70.3
BootBERT _{small}	74.9 \pm 3.4	-76.6 \pm 10.2	82.2	65.6

Table 1: The overall average scores for the four evaluation suites: (Super)GLUE, MSGS, BLiMP and BLiMP supplement. The (Super)GLUE and MSGS columns show the mean and standard deviation statistics across multiple runs. The best results for each track are typeset in bold. For a more complete view, the full distribution of the MSGS results is plotted in Figure 4.

(Matthews, 1975, which is renamed to the Linguistic Bias Score (LBS) in MSGS).

3.3 Results

The overall averaged results for all four evaluation suits are given in Table 1. Apart from evaluating masked autoencoders trained with latent bootstrapping (BootBERTs), as described in Section 2, we evaluate the three baseline language models provided by the organizers of BabyLM challenge: decoder-only OPT (Zhang et al., 2022), encoder-decoder T5 (Raffel et al., 2020) and encoder-only RoBERTa language models (Liu et al., 2019). As we base our models on the LTG-BERT architecture (Samuel et al., 2023a), we follow recommendations of the authors and also pretrain LTG-BERTs to get a strong and comparable baseline.

In addition to the averaged results, we also provide fine-grained (Super)GLUE scores in Table 2 and a visualization of the full distribution of MSGS scores in Figure 4 and in Appendix D (given the high variation of the aggregated MSGS results). The tables contain the mean and standard deviation statistics over 5 (respectively 15) runs. More details about the BLiMP and BLiMP supplement scores are given in Appendix A.

Model	BoolQ	CoLA	MNLI _m	MNLI _{mm}	MRPC	MultiRC	QNLI	QQP	RTE	SST-2	WSC	All
STRICT (100M words)												
OPT _{125m}	66.4 \pm 0.7	74.9 \pm 0.6	75.7 \pm 0.3	77.0 \pm 0.3	81.9 \pm 0.7	61.5 \pm 0.8	82.8 \pm 0.8	84.3 \pm 0.1	58.6 \pm 2.9	87.7 \pm 0.7	52.3 \pm 12.5	73.0 \pm 3.9
RoBERTa _{base}	67.7 \pm 0.7	75.6 \pm 0.3	77.4 \pm 0.4	78.3 \pm 0.3	84.0 \pm 0.5	64.3 \pm 0.5	83.6 \pm 0.2	85.5 \pm 0.2	50.7 \pm 1.5	88.3 \pm 0.6	61.4 \pm 0.0	74.3 \pm 0.6
T5 _{base}	67.7 \pm 1.5	76.7 \pm 0.9	77.9 \pm 0.3	78.7 \pm 0.3	85.2 \pm 1.1	65.7 \pm 0.8	84.7 \pm 0.9	86.2 \pm 0.1	55.4 \pm 2.2	89.0 \pm 0.8	61.0 \pm 1.1	75.3 \pm 1.1
LTG-BERT _{base}	68.1 \pm 0.4	82.8 \pm 0.4	83.4 \pm 0.3	83.1 \pm 0.2	84.3 \pm 0.7	71.2 \pm 0.9	89.3 \pm 0.3	87.9 \pm 0.2	55.2 \pm 2.7	91.9 \pm 0.6	58.6 \pm 3.5	77.8 \pm 1.4
BootBERT _{base}	72.4 \pm 1.2	81.6 \pm 0.6	84.7 \pm 0.3	84.7 \pm 0.3	89.1 \pm 0.3	70.7 \pm 1.2	91.2 \pm 0.4	88.1 \pm 0.1	57.2 \pm 3.5	91.8 \pm 0.8	60.2 \pm 2.7	79.2 \pm 1.5
STRICT-SMALL (10M words)												
OPT _{125m}	66.2 \pm 1.5	69.0 \pm 0.5	69.5 \pm 0.2	71.0 \pm 0.5	80.0 \pm 1.8	56.5 \pm 2.0	71.5 \pm 0.7	80.3 \pm 0.3	51.3 \pm 2.1	85.4 \pm 0.9	50.8 \pm 10.3	68.3 \pm 3.3
RoBERTa _{base}	65.8 \pm 2.9	70.4 \pm 0.4	72.5 \pm 0.4	74.4 \pm 0.3	82.2 \pm 0.4	61.2 \pm 1.5	80.3 \pm 0.7	83.5 \pm 0.2	56.8 \pm 5.5	85.6 \pm 0.3	61.7 \pm 0.5	72.2 \pm 1.9
T5 _{base}	63.4 \pm 1.6	69.4 \pm 0.1	57.3 \pm 0.8	58.6 \pm 1.1	81.4 \pm 0.6	48.4 \pm 1.4	64.3 \pm 0.9	76.8 \pm 0.3	52.7 \pm 2.4	79.4 \pm 1.0	60.0 \pm 2.2	64.7 \pm 1.3
LTG-BERT _{small}	64.8 \pm 2.1	77.6 \pm 0.8	78.0 \pm 0.2	78.8 \pm 0.4	82.3 \pm 0.4	64.1 \pm 0.3	85.0 \pm 0.2	85.8 \pm 0.2	53.7 \pm 4.1	88.8 \pm 0.8	60.5 \pm 1.0	74.5 \pm 1.5
BootBERT _{small}	67.6 \pm 2.7	75.3 \pm 1.4	79.2 \pm 0.3	80.0 \pm 0.2	83.2 \pm 1.5	65.2 \pm 0.9	86.2 \pm 0.4	86.6 \pm 0.1	54.7 \pm 5.3	88.2 \pm 0.8	57.3 \pm 9.2	74.9 \pm 3.4

Table 2: The BabyLM-flavored (Super)GLUE results of language models in the STRICT track and STRICT-SMALL track. We present the mean and standard deviation statistics over 5 finetuning runs (initialized with different random seeds) and boldface the best mean-result.

4 Discussion

LTG-BERT performance. Our results confirm the findings by Samuel et al. (2023a) who introduced the improved language modeling architecture called LTG-BERT. These models perform drastically better than the OPT, RoBERTa and T5 baselines pretrained on the same low-resource BabyLM corpus; the performance is improved across all evaluation suites – GLUE, MSGS, BLiMP as well as the BLiMP supplemental data – and across both STRICT and STRICT-SMALL tracks. LTG-BERT has also been used as the backbone of recent Norwegian language models trained on large amounts of data (Samuel et al., 2023b), which demonstrates that LM methods developed for efficient training are also beneficial for large-scale training.

Self-supervised learning. When we compare BootBERT to the LTG-BERT baseline, we can see that the latent bootstrapping approach leads to a substantially better performance when finetuned on (Super)GLUE in the STRICT track and to a slightly better performance in the STRICT-SMALL track. Specifically on the biggest and arguably most robust GLUE task, MNLI, the accuracy is better by 1.3/1.6 percentage points in the STRICT track and by 1.2/1.2pp in the STRICT-SMALL track. The overall average (Super)GLUE score is better by 1.4pp and by 0.4pp, respectively. This shows that language models pretrained with this approach are a good option for downstream tasks.

The ability of linguistic generalization, as measured by the linguistic bias scores in MSGS, is

substantially worse in BootBERT than in the LTG-BERT baseline, as evident from Figure 4. A more detailed analysis in Appendix D reveals that this holds for both BabyLM tracks – but the difference is mainly due to the fact that LTG-BERT reliably prefers the linguistic feature ‘is the main verb in “ing” form?’, other tests are relatively similar for both types of models. It is unclear what part of latent bootstrapping causes this difference.

The results on the BLiMP-based benchmarks are mixed but overall worse when comparing BootBERT with the LTG-BERT baseline. This is possibly because of the utilization of two conflicting training objectives in BootBERT – intuitively, pure language-modeling-based training should have an advantage on benchmarks that rely on sentence likelihood.

In conclusion, these low-resource experiments suggest that **the advantage of latent bootstrapping for natural language is not as great as the advantage that has been previously demonstrated for computer vision.** We believe that this is because the atomic units of text, subword tokens, can provide much more semantically rich signal when compared to the atomic units of images, pixels. Thus there is not a large need for bootstrapping a rich signal from a teacher; instead, the standard language modeling comes with a training objective that is simple and provides enough signal, while suffering from issues like representation collapse.

The shared task results. The official DynaBench results for BabyLM can be found in Ap-

pendix E. Our system ranks high when evaluated on GLUE (first and second place) and on BLiMP (second and first place) in the STRICT and STRICT-SMALL tracks, respectively. As discussed earlier, BootBERT strongly prefers the surface features over the linguistic features and thus places low on the MSGS benchmark (third and last place), which also hurts the overall ranking of our system (third an seventh). Note however that this evaluation is not using a proper train/development/test split and it does not account for high variation of some metrics (MSGS in particular), which is why we have used an alternative evaluation in the rest of this paper.

Computational cost of latent bootstrapping. It is important to note that latent bootstrapping comes with an increased computational cost because of an additional forward pass through the mean teacher; which roughly equates to a 50% increase in pre-training time. Thus, it should be carefully considered whether the potential benefits of bootstrapping are worth this cost. That being said, this method does not bear any additional cost during finetuning nor inference, which might justify it in some cases.

5 Related work

Self-supervised learning. Our work is greatly inspired by the ‘bootstrap your own latent’ approach (BYOL; Grill et al., 2020), which introduced the bootstrapping feedback loop between a student and a mean teacher network. BYOL by itself can be considered an example of contrastive learning (Hjelm et al., 2019; van den Oord et al., 2019; Chen et al., 2020; He et al., 2020) without negative instances. Another important aspect of BYOL is the usage of a ‘mean teacher’, a slow-moving average of a student network, which is a term coined by Tarvainen and Valpola (2017).

Many methods of visual representation learning adopted the bootstrapping approach and further improved its parts (Chen and He, 2021; Zbontar et al., 2021; Bardes et al., 2022; He et al., 2022). In particular, our work bears similarities with the recently introduces ‘image-based joint-embedding predictive architecture’ (I-JEPA; Assran et al., 2023), which also trains a masked autoencoder student network to predict the contextualized embeddings of an unmasked mean teacher. While mostly used for the image domain, *data2vec* method showed that latent bootstrapping can also be successfully applied to text (Baevski et al., 2022).

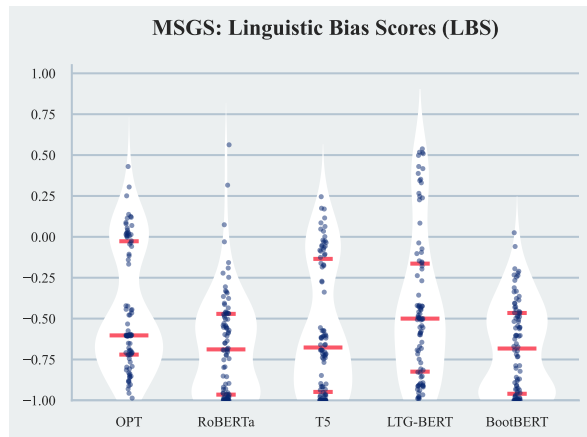


Figure 4: The Linguistic Bias Scores (LBS) of language models pretrained on the STRICT dataset. These plots show the distribution of the LBS scores across 15 evaluation runs (3 learning rates \times 5 random seeds) for each of the 6 non-ambiguous test datasets (90 values in total for each model). The red horizontal lines highlight the first, second (median) and third quartile. The overall negative scores show that none of the tested models prefers linguistic features over the surface ones.

Efficient language modeling. The necessity of pretraining modern language models on large corpora were questioned in CamemBERT (Martin et al., 2020) and the effect of corpus size has been then thoroughly studied in Micheli et al. (2020), Zhang et al. (2021) as well as in Hoffmann et al. (2022). Samuel et al. (2023a) introduced the LTG-BERT – an improved language model optimized for pretraining on a low-resource corpus. They showed that a well-tuned language model can match the performance of BERT even when it is pretrained only on a small 100-million-word British National Corpus (BNC). We base our approach on this model due to the apparent similarity of the BabyLM training corpus to BNC.

6 Conclusion

In this paper, we presented a masked autoencoder language model trained with latent bootstrapping, an alternative self-supervised learning method. We showed that when pretrained on a low-resource corpus, the results of this method are varied – compared to a masked language modeling baseline, the performance is clearly better on (Super)GLUE, but worse on MSGS and mixed on BLiMP. We believe that it makes a promising alternative to traditional language modeling methods, but its reliable and effective utilization requires future work.

Acknowledgements

This paper would not be possible without the endless support and incredibly useful feedback from Andrey Kutuzov, Erik Velldal and Lilja Øvrelid from the Language Technology Group at the University of Oslo.

The efforts described in the current paper were funded by the HPLT project (High Performance Language Technologies; coordinated by Charles University). The computations were performed on resources provided through Sigma2 – the national research infrastructure provider for High-Performance Computing and large-scale data storage in Norway.

References

- Ahmed Abdelali, Francisco Guzman, Hassan Sajjad, and Stephan Vogel. 2014. The AMARA corpus: Building parallel language resources for the educational domain. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. 2023. Self-supervised learning from images with a joint-embedding predictive architecture. *arXiv preprint arXiv:2301.08243*.
- Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatao Gu, and Michael Auli. 2022. [data2vec: A general framework for self-supervised learning in speech, vision and language](#). In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 1298–1312. PMLR.
- Roy Bar-Haim, Ido Dagan, Bill Dolan, Lisa Ferro, and Danilo Giampiccolo. 2006. The second pascal recognising textual entailment challenge. *Proceedings of the Second PASCAL Challenges Workshop on Recognising Textual Entailment*.
- Adrien Bardes, Jean Ponce, and Yann LeCun. 2022. [VICReg: Variance-invariance-covariance regularization for self-supervised learning](#). In *International Conference on Learning Representations*.
- Yonatan Belinkov. 2022. [Probing Classifiers: Promises, Shortcomings, and Advances](#). *Computational Linguistics*, 48(1):207–219.
- Luisa Bentivogli, Ido Dagan, Hoa Trang Dang, Danilo Giampiccolo, and Bernardo Magnini. 2009. The fifth pascal recognizing textual entailment challenge. In *In Proc Text Analysis Conference (TAC'09)*.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning, ICML'20*. JMLR.org.
- Xinlei Chen and Kaiming He. 2020. Exploring simple siamese representation learning. *arXiv preprint arXiv:2011.10566*.
- Xinlei Chen and Kaiming He. 2021. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15750–15758.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [BoolQ: Exploring the surprising difficulty of natural yes/no questions](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2924–2936, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2006. The pascal recognising textual entailment challenge. In *Machine Learning Challenges. Evaluating Predictive Uncertainty, Visual Object Classification, and Recognising Textual Entailment*, pages 177–190, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Allyson Ettinger. 2020. [What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models](#). *Transactions of the Association for Computational Linguistics*, 8:34–48.
- Leo Gao, Jonathan Tow, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Kyle McDonell, Niklas Muennighoff, Jason Phang, Laria Reynolds, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2021. [A framework for few-shot language model evaluation](#).
- Martin Gerlach and Francesc Font-Clos. 2018. [A standardized Project Gutenberg corpus for statistical analysis of natural language and quantitative linguistics](#). *Computing Research Repository*, arXiv:1812.08092.

- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. 2007. [The third PASCAL recognizing textual entailment challenge](#). In *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, pages 1–9, Prague. Association for Computational Linguistics.
- Ross Girshick. 2015. [Fast r-cnn](#). In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)*, ICCV '15, page 1440–1448, USA. IEEE Computer Society.
- Thamme Gowda and Jonathan May. 2020. [Finding the optimal vocabulary size for neural machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3955–3964, Online. Association for Computational Linguistics.
- Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. 2020. Bootstrap your own latent a new approach to self-supervised learning. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS'20*, Red Hook, NY, USA. Curran Associates Inc.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017a. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17*, page 1321–1330. JMLR.org.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017b. On calibration of modern neural networks.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16000–16009.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. [Momentum contrast for unsupervised visual representation learning](#). In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9726–9735.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [Deberta: Decoding-enhanced bert with disentangled attention](#). In *International Conference on Learning Representations*.
- Felix Hill, Antoine Bordes, Sumit Chopra, and Jason Weston. 2016. [The Goldilocks principle: Reading children's books with explicit memory representations](#).
- R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. 2019. [Learning deep representations by mutual information estimation and maximization](#). In *International Conference on Learning Representations*.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre. 2022. [Training compute-optimal large language models](#).
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. [SpanBERT: Improving pre-training by representing and predicting spans](#). *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Daniel Khashabi, Snigdha Chaturvedi, Michael Roth, Shyam Upadhyay, and Dan Roth. 2018. [Looking beyond the surface: A challenge set for reading comprehension over multiple sentences](#). 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 252–262, New Orleans, Louisiana. Association for Computational Linguistics.
- Yann LeCun. 2022. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Review*, 62.
- Hector J. Levesque, Ernest Davis, and Leora Morgenstern. 2012. The winograd schema challenge. In *Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning, KR'12*, page 552–561. AAAI Press.
- Pierre Lison and Jörg Tiedemann. 2016. [OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 923–929, Portorož, Slovenia. European Language Resources Association (ELRA).
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Brian MacWhinney. 2000. *The CHILDES project: The database*, volume 2. Psychology Press.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric de la Clergerie, Djamé Seddah, and Benoît Sagot. 2020. [CamemBERT: a tasty French language model](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7203–7219, Online. Association for Computational Linguistics.
- B.W. Matthews. 1975. [Comparison of the predicted and observed secondary structure of t4 phage lysozyme](#). *Biochimica et Biophysica Acta (BBA) - Protein Structure*, 405(2):442–451.

- Yu Meng, Jitin Krishnan, Sinong Wang, Qifan Wang, Yuning Mao, Han Fang, Marjan Ghazvininejad, Jiawei Han, and Luke Zettlemoyer. 2023. [Representation deficiency in masked language modeling](#).
- Vincent Micheli, Martin d’Hoffschmidt, and François Fleuret. 2020. [On the importance of pre-training data volume for compact language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7853–7858, Online. Association for Computational Linguistics.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. [Improving language understanding by generative pre-training](#).
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. 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.
- David Samuel, Andrey Kutuzov, Lilja Øvrelid, and Erik Velldal. 2023a. [Trained on 100 million words and still in shape: BERT meets British National Corpus](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1954–1974, Dubrovnik, Croatia. Association for Computational Linguistics.
- David Samuel, Andrey Kutuzov, Samia Touileb, Erik Velldal, Lilja Øvrelid, Egil Rønningstad, Elina Sigdel, and Anna Palatkina. 2023b. [NorBench – a benchmark for Norwegian language models](#). In *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 618–633, Tórshavn, Faroe Islands. University of Tartu Library.
- Noam Shazeer. 2020. [GLU variants improve transformer](#). *CoRR*, abs/2002.05202.
- Sam Shleifer and Myle Ott. 2022. [Normformer: Improved transformer pretraining with extra normalization](#).
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Andreas Stolcke, Klaus Ries, Noah Coccaro, Elizabeth Shriberg, Rebecca Bates, Daniel Jurafsky, Paul Taylor, Rachel Martin, Marie Meteer, and Carol Van Ess-Dykema. 2000. Dialogue act modeling for automatic tagging and recognition of conversational speech. *Computational Linguistics*, 26(3):339–371.
- Antti Tarvainen and Harri Valpola. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 1195–1204, Red Hook, NY, USA. Curran Associates Inc.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Sam Bowman, Dipanjan Das, and Ellie Pavlick. 2019. [What do you learn from context? probing for sentence structure in contextualized word representations](#). In *International Conference on Learning Representations*.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2019. [Representation learning with contrastive predictive coding](#).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Alex Wang and Kyunghyun Cho. 2019. [BERT has a mouth, and it must speak: BERT as a Markov random field language model](#). In *Proceedings of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation*, pages 30–36, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- 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.
- Alex Warstadt, Leshem Choshen, Aaron Mueller, Adina Williams, Ethan Wilcox, and Chengxu Zhuang. 2023a. [Call for papers – the babyLM challenge: Sample-efficient pretraining on a developmentally](#)

- plausible corpus. *Computing Research Repository*, arXiv:2301.11796.
- Alex Warstadt, Aaron Mueller, Leshem Choshen, Ethan Gotlieb Wilcox, Chengxu Zhuang, Juan Ciro, Rafael Mosquera, Adina Williams, Bhargavi Paranjabe, Tal Linzen, and Ryan Cotterell. 2023b. 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 Mohananeey, 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, Amanpreet Singh, and Samuel R. Bowman. 2019. **Neural network acceptability judgments**. *Transactions of the Association for Computational Linguistics*, 7:625–641.
- 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.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. **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, New Orleans, Louisiana. Association for Computational Linguistics.
- C. Wilson. 2006. Learning phonology with substantive bias: an experimental and computational study of velar palatalization. *Cogn Sci*, 30(5):945–982.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. **Google’s neural machine translation system: Bridging the gap between human and machine translation**.
- Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stephane Deny. 2021. **Barlow twins: Self-supervised learning via redundancy reduction**. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 12310–12320. PMLR.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. 2022. **Opt: Open pre-trained transformer language models**.
- 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.

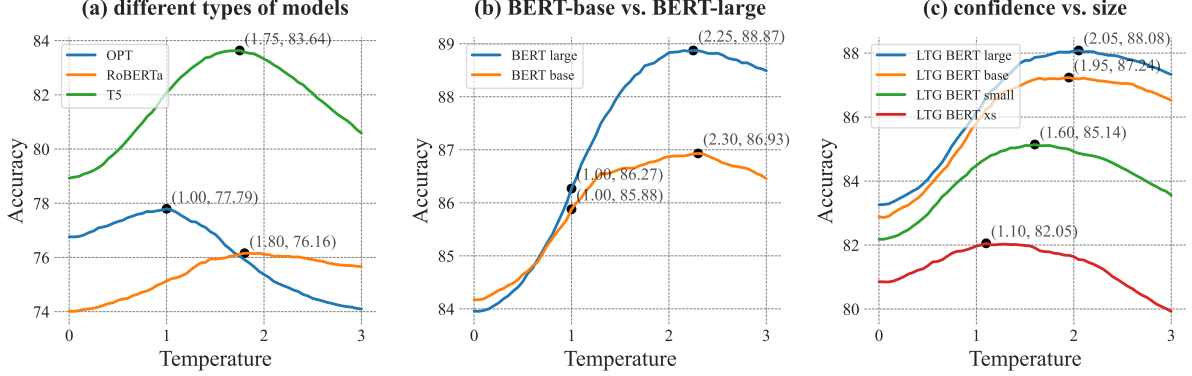


Figure 5: These plots show the BLiMP ‘confidence profiles’ of several language models – the influence of temperature scaling on the average BLiMP accuracy. (a) Models trained by different training objectives show different confidence profiles, judging their linguistic knowledge from BLiMP accuracy can be misleading. Here, we compare the three baseline from the BabyLM challenge trained on the STRICT track. (b) The linguistic knowledge of BERT_{base} and BERT_{large} appears comparable when judging from performance at temperature 1, but the potential of the larger model is much greater. (c) We train four sizes of LTG-BERT on the STRICT track and plot their confidence profiles. Larger models tend to be more confident and, therefore, measuring them at temperature 1 is more misleading.

A The Effect of temperature scaling on BLiMP

Our preliminary experiments with calibrating language models via temperature scaling (Guo et al., 2017b) revealed that the BLiMP scores are hugely dependent on the scalar temperature parameter – when these are calculated with the standard method by (Salazar et al., 2020). This single temperature value can increase the accuracy on some BLiMP subtasks by more than 10% (Figure 6), which challenges the usage of BLiMP as an appropriate evaluation tool. It is especially problematic when comparing different types of language models (Figure 5a) and different sizes of language models (Figure 5b,c).

Background. To better understand this problem, this section describes how are the BLiMP scores traditionally computed for masked language models. These models can estimate $P(s_t | s_{\setminus t})$ – the likelihood of a token s_t given its bidirectional context $s_{\setminus t} = (s_i | i \neq t)$. This probability distribution P is given by a softmax transformation of the output logits z , where τ is temperature:

$$P_i = \frac{\exp(z_i/\tau)}{\sum_k \exp(z_k/\tau)}.$$

Large temperature yields more even distribution and low temperature gives more ‘peaky’ distribution.

Salazar et al. (2020) proposed to use these probability estimates (with $\tau = 1$) to infer a *score* for each BLiMP sentence, with a higher *score* corresponding to a more likely sentence. Then, the BLiMP accuracy measures how many times is the score of a grammatically correct sentence greater than the score of an incorrect sentence. Specifically, we use the *pseudo-log-likelihood score* (PPL) by Wang and Cho (2019). The PPL score of a sentence s is defined as:

$$\text{PPL}(s) = \sum_{t=1}^N \log P(s_t | s_{\setminus t}).$$

Proposed solution. BLiMP should measure the linguistic knowledge of language models and we believe that this metric should be independent of the prediction confidence of these models. Formally speaking, the BLiMP score should be invariant to temperature scaling. Therefore, we propose to use the maximal average accuracy across all possible temperature values – instead of simply using the average accuracy at temperature equal to 1. As apparent from Figure 5b, such formulation can better reflect the difference of linguistic knowledge found in BERT_{base} and BERT_{large}. There, the accuracy measured at temperature 1

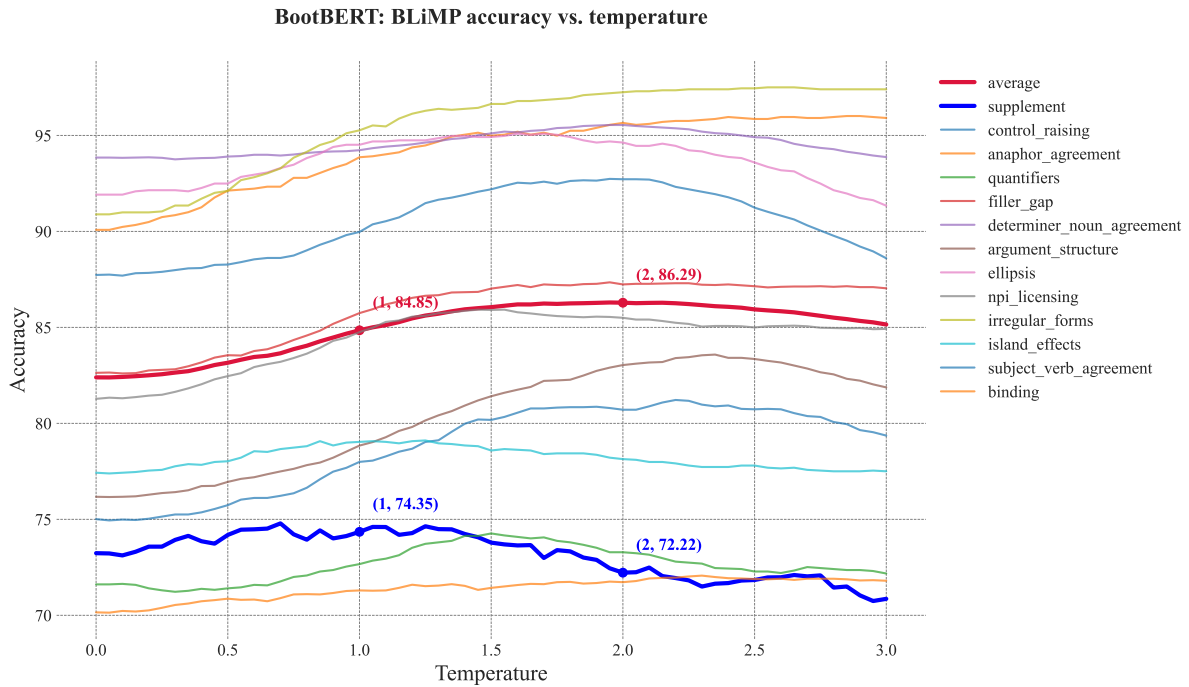


Figure 6: The confidence profile of our proposed $\text{BootBERT}_{\text{base}}$ model pretrained on the STRICT track. Apart from the average BLiMP accuracy (in red) and the average BLiMP supplement accuracy (in blue), this plot shows fine-grained BLiMP accuracies on all subtasks.

is at odds with other measures that show substantially better linguistic knowledge of $\text{BERT}_{\text{large}}$ (Devlin et al., 2019; Tenney et al., 2019; Ettinger, 2020).

Note that our approach bears only a negligible compute cost because the temperature modification is done ex-post, i.e., it does not require any additional passes through the language model.

Using one temperature for all subtasks does not account for the severe difference between the accuracy scores on these tasks (Figure 6), but it is a simple solution that also allows us to evaluate models on a held-out set, such as the BLiMP supplement. We believe that a scoring function that is (i) unified, (ii) invariant to temperature and (iii) fair to all subtasks, is an interesting future work.

B Data preprocessing

The pretraining datasets for the STRICT and STRICT-SMALL tracks are a mix 10 different corpora, as shown in Table 3. We applied light preprocessing and normalization to these corpora in order to cast them into a unified format. In particular, we applied these modifications:

- **CHILDES:** We capitalize the first letter of each line, normalize punctuation with whitespaces (essentially detokenization) and put every line between double quotes (as directed speech).
- **British National Corpus:** Capitalization, normalization and double quotes.
- **Children’s Book Test:** This corpus contains some remnants of the Penn Tree format where, for example, -LRB- and -RRB- tokens are used instead of ‘(’ and ‘)’. We normalize all unnatural symbols and whitespaces.
- **Children’s Stories Text Corpus:** We try to conserve the formatting with a special [TAB] symbol and apply whitespace normalization.
- **Standardized Project Gutenberg Corpus:** The text file is aligned into blocks by inserting a newline symbol after at most 70 characters, which ruins the sentence structure. We restore the original paragraphs by removing these additional newline symbols and apply whitespace normalization.

- **OpenSubtitles:** Some lines arbitrarily start with a dash symbol, which we remove. Then whitespace normalization is applied and every line is cast a directed speech with double quotes.
- **QED:** This corpus contains some incorrectly parsed HTML symbols, which we tried to clean up with some simple heuristics. The whitespace normalization is applied and every line is cast as directed speech with double quotes.
- **Wikipedia:** This dataset also needed to be cleaned of incorrectly parsed Wikipedia tags and hyperlinks. Whitespace normalization is applied.
- **Simple Wikipedia:** Heuristic HTML clean-up and whitespace normalization.
- **Switchboard:** The same as OpenSubtitles: removed leading dashes, whitespaces normalization and added double quotes.

Note that the preprocessed corpora and the preprocessing scripts are released alongside the training scripts.

Dataset	Domain	# Words		
		STRICT-SMALL	STRICT	Proportion
CHILDES (MacWhinney, 2000)	Child-directed speech	0.44M	4.21M	5%
British National Corpus (BNC), ¹ dialogue portion	Dialogue	0.86M	8.16M	8%
Children’s Book Test (Hill et al., 2016)	Children’s books	0.57M	5.55M	6%
Children’s Stories Text Corpus ²	Children’s books	0.34M	3.22M	3%
Standardized Project Gutenberg Corpus (Gerlach and Font-Clos, 2018)	Written English	0.99M	9.46M	10%
OpenSubtitles (Lison and Tiedemann, 2016)	Movie subtitles	3.09M	31.28M	31%
QCRI Educational Domain Corpus (QED; Abdelali et al., 2014)	Educational video subtitles	1.04M	10.24M	11%
Wikipedia ³	Wikipedia (English)	0.99M	10.08M	10%
Simple Wikipedia ⁴	Wikipedia (Simple English)	1.52M	14.66M	15%
Switchboard Dialog Act Corpus (Stolcke et al., 2000)	Dialogue	0.12M	1.18M	1%
<i>Total</i>	–	9.96M	98.04M	100%

Table 3: The contents of datasets for the the STRICT and STRICT-SMALL tracks; the table is taken from Warstadt et al. (2023b). ¹<http://www.natcorp.ox.ac.uk> ²<https://www.kaggle.com/datasets/edenbd/children-stories-text-corpus> ³<https://dumps.wikimedia.org/enwiki/20221220/> ⁴<https://dumps.wikimedia.org/simplewiki/20221201/>

C Implementation details

In order to reduce training time, pre-training is parallelized over multiple GPUs with the global batch size of 4 096. The number of GPUs used depends on the size of pre-trained language models, ranging from 32 to 128 AMD Instinct MI250X GPUs, each with 64GB memory. The amount of training steps is 62 500, reducing the training budget of the original BERT model by 50%. Unlike the BERT and LTG-BERT training recipe, we use the same sequence length, 256, throughout the whole training. This decision is necessary for keeping a reasonable exponential moving average of the parameters (it could be corrupted when switching to a longer sequence length in the middle of training).

The implementation of latent bootstrapping mainly follows I-JEPA (Assran et al., 2023). We also adopt their usage of a linearly increasing schedule of the EMA decay hyperparameter τ and a cosine schedule of weight decay.

The hyperparameters for pretraining are given in Table 5. Table 6 shows the finetuning hyperparameters.

D Finegrained MSGS scores

This section shows the full score distribution over all MSGS subtasks, including the control subtasks. This gives a better view on the behavior of different language models than the aggregated scores in [Figure 4](#) and [Table 1](#).

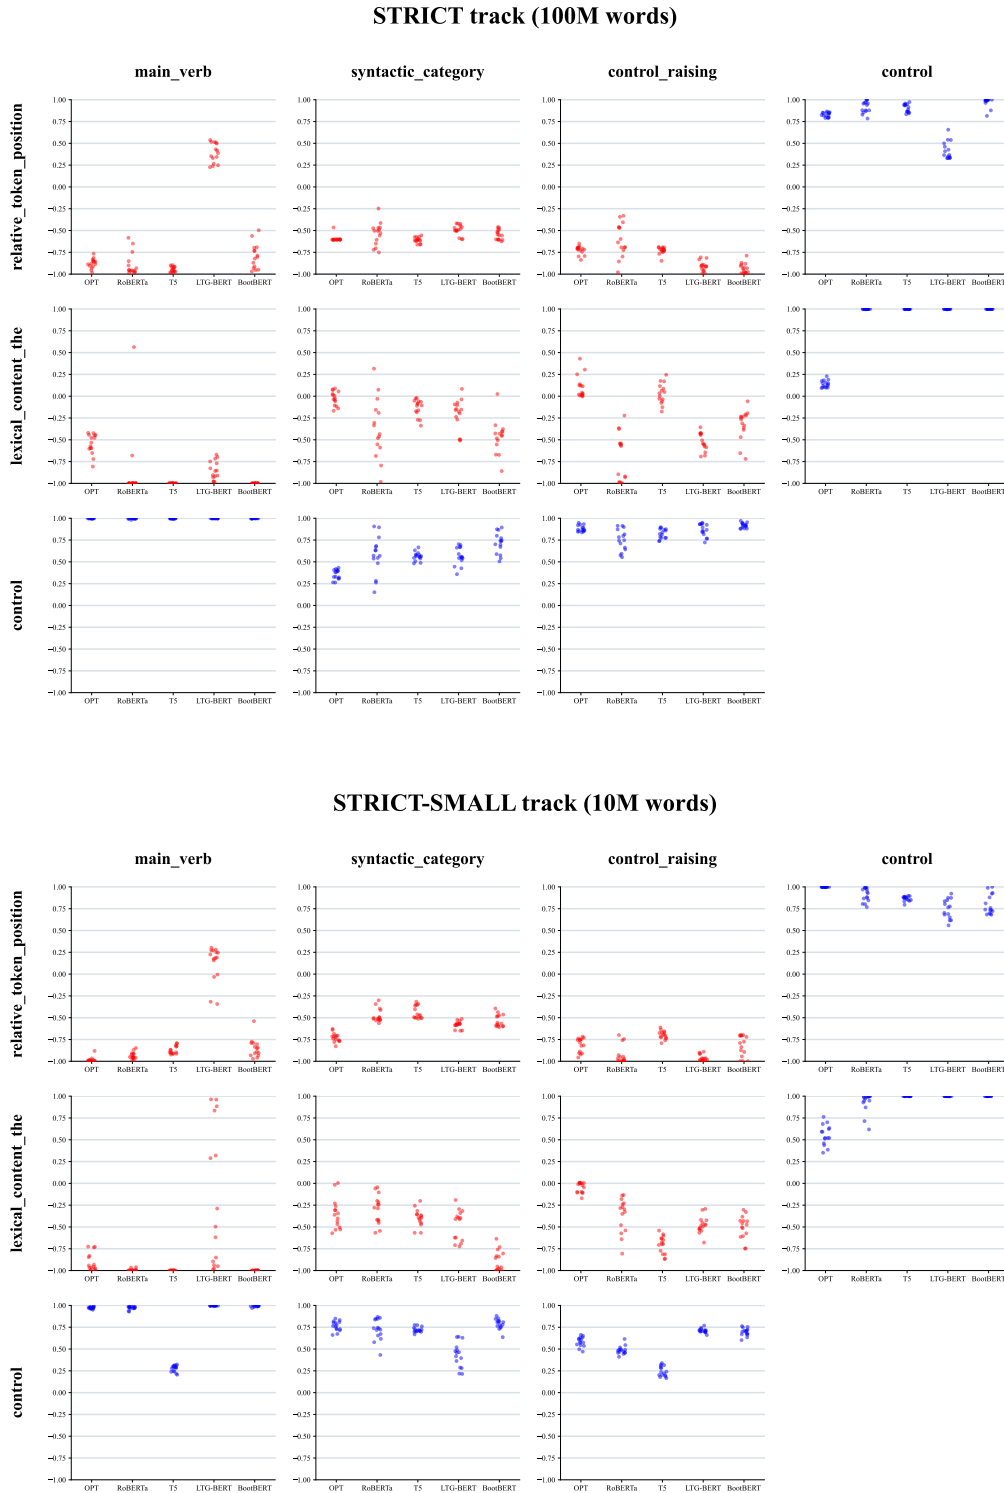


Figure 7: The MSGS linguistic bias scores of the control tasks (in blue) and non-control disambiguated tasks (in red). Values close to 1 indicate preference of linguistic explanations (columns) while values close to -1 indicate preference of surface explanations.

E The official BabyLM results from DynaBench

This section shows the official results for the BabyLM challenge as published on the DynaBench website.⁵ We show the top 9 submissions (the official ones delivered on time) for the STRICT and STRICT-SMALL tracks with the aggregated scores in Table 4.

STRICT track (100M words)					STRICT-SMALL track (10M words)				
Model	BLiMP	GLUE	MSGS	Average	Model	BLiMP	GLUE	MSGS	Average
BootBERT	#2 82.2	#1 78.5	#3 27.7	#3 70.2	BootBERT	#1 75.9	#2 71.7	#9 -9.7	#7 57.5
ELC-BERT	82.8	78.3	47.2	74.3	ELC-BERT	75.8	73.7	29.4	65.9
Contextualizer	79.0	72.9	58.0	73.0	MLSM	72.4	70.6	17.2	60.8
MSLM	76.2	73.5	21.4	64.4	Contextualizer	74.3	69.6	12.7	60.5
Bad babies	77.0	67.2	23.4	63.4	Baby Llama	69.8	67.6	24.7	60.1
CogMemLM	72.8	72.2	-0.1	58.0	Too Much Information	75.7	70.9	3.9	59.9
Pre-training LLMs	71.6	69.8	-3.8	56.0	McGill	72.4	69.3	5.2	58.0
BabyStories	73.9	59.1	0.2	54.7	CLIMB	71.8	65.6	9.7	57.5
AB-RoBERTa	68.3	64.1	-11.8	51.0	William’s college GPT2	70.9	64.8	9.9	56.9

Table 4: The DynaBench scores of the BabyLM challenge (Warstadt et al., 2023a), the table shows the top 9 submissions in the STRICT and STRICT-SMALL tracks. Higher scores are better, the best results in each evaluation suite are boldfaced.

F BabyLM subset of (Super)GLUE tasks

The BabyLM challenge involves slightly modified GLUE and SuperGLUE benchmarks. It uses only a subset of the subtasks, the datasets are filtered so that they do not contain out-of-vocabulary words, and it sometimes use non-standard metrics. We list all subtasks and their metrics below:

- **Boolean Questions** (BoolQ; Clark et al., 2019), a yes/no Q/A dataset evaluated with accuracy.
- **Corpus of Linguistic Acceptability** (CoLA; Warstadt et al., 2019) evaluated with accuracy (originally evaluated with the Matthews correlation coefficient (MCC; Matthews, 1975)).
- **The Multi-Genre Natural Language Inference Corpus** (MNLI; Williams et al., 2018). Its development set consists of two parts: *matched*, sampled from the same data source as the training set, and *mismatched*, which is sampled from a different domain. Both parts are evaluated with accuracy.
- **The Microsoft Research Paraphrase Corpus** (MRPC; Dolan and Brockett, 2005), evaluated with both F₁-score (originally also evaluated with accuracy).
- **Multi-Sentence Reading Comprehension** (MultiRC; Khashabi et al., 2018), a multiple choice question answering dataset, evaluated with accuracy (originally evaluated with the exact match accuracy (EM) and F₁-score (over all answer options)).
- **Question-answering Natural Language Inference** (QNLI) constructed from the Stanford Question Answering Dataset (SQuAD; Rajpurkar et al., 2016), evaluated with accuracy.
- **The Quora Question Pairs** (QQP),⁶ evaluated with F₁-score (originally evaluated with accuracy).
- **The Stanford Sentiment Treebank** (SST-2; Socher et al., 2013), evaluated with accuracy.
- **The Recognizing Textual Entailment datasets** (RTE; Dagan et al., 2006; Bar-Haim et al., 2006; Giampiccolo et al., 2007; Bentivogli et al., 2009), evaluated with accuracy.
- **Winograd Schema Challenge** (WSC; Levesque et al., 2012) evaluated with accuracy.

⁵<https://dynabench.org/babylm> (22 October, 2023)

⁶<https://quoradata.quora.com/First-Quora-Dataset-Release-Question-Pairs>

Hyperparameter	BootBERT _{small}	BootBERT _{base}
Number of parameters	30 395 776	127 744 768
Number of layers	12	12
Hidden size	384	768
FF intermediate size	1 024	2 048
Vocabulary size	4 096	16 384
Attention heads	6	12
β parameter	0.1	0.1
Encoder hidden dropout	0.1	0.1
Encoder attention dropout	0.1	0.1
Decoder hidden dropout	0.0	0.0
Decoder attention dropout	0.0	0.0
Training steps	62 500	62 500
Batch size	4 096	4 096
Sequence length	256	256
Warmup steps	1 000	1 000
Initial learning rate	0.007	0.005
Final learning rate	0.0007	0.0005
Learning rate scheduler	cosine	cosine
Initial weight decay	0.04	0.02
Final weight decay	0.4	0.2
Weight decay scheduler	cosine	cosine
Initial EMA decay τ	0.996	0.996
Final EMA decay τ	1.0	1.0
EMA decay scheduler	linear	linear
Layer norm ϵ	1e-7	1e-7
Optimizer	LAMB	LAMB
LAMB ϵ	1e-6	1e-6
LAMB β_1	0.9	0.9
LAMB β_2	0.98	0.98
Gradient clipping	2.0	2.0

Table 5: Pre-training hyperparameters for the small-sized BootBERT (trained on STRICT-SMALL and for the base-sized BootBERT (trained on the STRICT track).

Hyperparameter	BoolQ, MNLI		
	MRPC, MultiRC, QNLI	CoLA, RTE, WSC	MSGS
	QQP, SST-2		
Batch size	32	16	16
Number of epochs	10	10	5
Dropout	0.1	0.1	0.1
Warmup steps	10%	10%	10%
Peak learning rate	3e-5	3e-5	{1e-5, 2e-5, 3e-5}
Learning rate decay	linear	linear	linear
Weight decay	0.01	0.01	0.01
Optimizer	AdamW	AdamW	AdamW
Adam ϵ	1e-6	1e-6	1e-6
Adam β_1	0.9	0.9	0.9
Adam β_2	0.999	0.999	0.999

Table 6: Hyperparameters for fine-tuning the GLUE, SuperGLUE task and MSGS tasks. We use the same hyperparameters for all models, not performing any per-model hyperparameter search. These values are adopted from LTG-BERT (Samuel et al., 2023a) and MSGS (Warstadt et al., 2020b). For all models, we measure the statistics over 5 random seeds: 1234, 2345, 3456, 4567 and 5678.