

Repurposing Language Models into Embedding Models: Finding the Compute-Optimal Recipe

Albert Q. Jiang*
University of Cambridge

Alicja Ziarko*
IDEAS NCBR
University of Warsaw
IMPAN

Bartosz Piotrowski
IDEAS NCBR

Wenda Li
University of Edinburgh

Mateja Jamnik†
University of Cambridge

Piotr Miłoś†
IDEAS NCBR
University of Warsaw
IMPAN, deepsense.ai

Abstract

Text embeddings are essential for many tasks, such as document retrieval, clustering, and semantic similarity assessment. In this paper, we study how to contrastively train text embedding models in a compute-optimal fashion, given a suite of pre-trained decoder-only language models. Our innovation is an algorithm that produces optimal configurations of model sizes, data quantities, and fine-tuning methods for text-embedding models at different computational budget levels. The resulting recipe, which we obtain through extensive experiments, can be used by practitioners to make informed design choices for their embedding models. Specifically, our findings suggest that full fine-tuning and low-rank adaptation fine-tuning produce optimal models at lower and higher computational budgets respectively.

1 Introduction

Text embeddings are vector representations of text sequences [Devlin et al., 2019]. A desired property of text embeddings is that their distribution in the embedding space reflects the semantic relations of the embedded texts. This property benefits multiple practical tasks [Muenighoff et al., 2023b], such as document retrieval (where documents are ranked based on the distance between their representations and that of a query), document clustering, sentiment analysis, or textual similarity measurement.

In the era of large language models (LLMs) that are pre-trained on vast textual data, it is a natural idea to capitalise on the language representations learned by these models to achieve high-quality embedding models. A common and effective approach has been to initialise an embedding model from a pre-trained LLM and fine-tune it with a *contrastive loss* [Neelakantan et al., 2022, Izacard et al., 2022]. This phase is

*Equal contribution.

†Equal advising contribution.

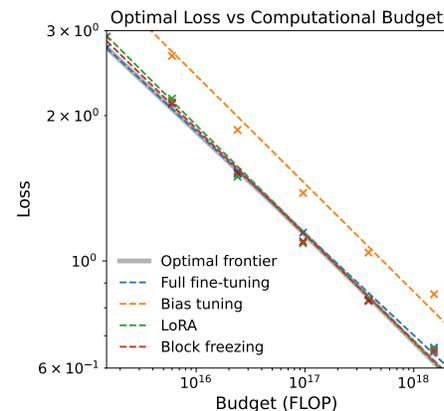


Figure 1: The optimal loss achieved using four different fine-tuning methods (full fine-tuning, only tuning the bias, low-rank adaptation, and freezing transformer blocks) at given budgets. The horizontal axis is the computational budget in floating point operations (FLOP) and the vertical axis is the contrastive loss. The X marks are datapoints and dotted lines are fitted linear trends for different methods. The solid black line is the “optimal frontier,” i.e., the optimal loss achievable with a fixed budget and the best method.

of crucial importance, as the quality of the embeddings extracted directly from a pre-trained LLM (e.g., by averaging hidden states of the last layer) is not sufficient for most tasks.

However, the best performing modern LLMs are usually massive in terms of the parameter counts (e.g., 175B parameters for GPT3 [Brown et al., 2020] and 340B for PaLM2 [Anil et al., 2023]), and are therefore difficult to train with limited resources. This motivates a practical question, that to the best of our knowledge has not been addressed systematically, and thus we address it in this paper: *what is the best embedding model one can train from a backbone decoder-only LLM with a fixed training compute budget?*

To answer this, we performed an extensive empirical study. We started by identifying design choices one can make when fine-tuning language models into embedding models, including: model size, data volume, parameter-efficient fine-tuning technique, and the hyperparameters of the chosen technique. Then, we performed a grid search over the pre-defined design choices, and found the best performing configuration under each computational budget. Using these findings, we established scaling laws for contrastive fine-tuning, which enabled us to build an algorithm that produces a general recipe for efficient fine-tuning of a decoder-only LLM that obtains a high-quality embedding model.

Contribution We comprehensively experimented with contrastively fine-tuning language models into text embedding models. We analysed how the choice of model sizes, data quantities, and fine-tuning methods affect the performance in the resource-constrained training regime. We compiled the findings into an algorithm that, given a fixed computational budget, predicts the optimal network architecture, data quantity, and parameter-efficient fine-tuning hyperparameters. We open-source the code to train and evaluate our models at: <https://github.com/SeqDM/Efficient-Embeddings>.

2 Related work

Embedding models Training neural networks to represent text as continuous vectors was popularised by word2vec [Mikolov et al., 2013], which produced semantically meaningful embeddings of *words*. BERT [Devlin et al., 2019] and the further contrastively trained SimCSE [Gao et al., 2021c] quickly established encoder-only transformers as the go-to architecture for embedding *text*.

More recently, decoder-only transformers have become increasingly more powerful and efficient at the same time [Brown et al., 2020, Touvron et al., 2023, Jiang et al., 2023]. Building an embedding model based on them utilises the knowledge from their pretraining, and thus it is a natural step. Neelakantan et al. [2022] set a successful precedent in this paradigm by initialising embedder training with decoder-only GPT models.

Notably, the current open state-of-the-art (the top open model on the MTEB [Muennighoff et al., 2023b] leaderboard, as of 18 May 2024), SFR-Embedding-Mistral [Meng et al., 2024], is also fine-tuned from the decoder-only Mistral 7B [Jiang et al., 2023].

Benchmarks for embedding models Embedding models are versatile in their applications. Therefore, a broad and diverse benchmark to provide a robust measure of their performance is called for. The first such benchmark was BEIR introduced by Thakur et al. [2021]. It has nine different information retrieval tasks (e.g., duplicate-question retrieval or citation-prediction) on 18 datasets. Recently, Muennighoff et al. [2023b] introduced MTEB (Massive Text Embedding Benchmark), which is substantially larger than BEIR and, in addition to information retrieval, has seven other types of tasks (e.g., clustering, summarisation, and pair classification) using 58 datasets and covering 112 languages. The authors also provide a leaderboard that currently contains more than 220 entries.

Scaling laws Scaling laws predict the performance of models at larger scale (e.g., more parameters, more data) from experiments at smaller scales. Kaplan et al. [2020] and Hoffmann et al. [2022] used empirically fitted laws to predict the performance of neural language models at different model sizes and data scales, making it possible to calculate the optimal configuration before training. Muennighoff et al. [2023a] inspected model training in data-constrained, multi-epoch regime. Frantar et al. [2023] looked at scaling laws for language models in the context of weight sparsity. Zhang et al. [2024] investigates the scaling laws for model fine-tuning, including parameter efficient methods: LoRA [Hu et al., 2022] and prompt tuning [Lester et al., 2021]. They show that the scaling laws and the best fine-tuning method are task dependent. Biderman et al. [2024] investigates the properties of fine-tuning models with LoRA in detail, showing that LoRA underperforms full fine-tuning when fine-tuning the models on mathematical or code datasets.

The single most related investigation is a concurrent work Fang et al. [2024], where the scaling laws for encoder models for retrieval are investigated. There are significant differences in the settings we consider. Our focus is on investigating the process of fine-tuning decoder-only models for good quality embeddings. Moreover, our main goal is to find which strategy for achieving good embeddings is optimal in a budget-restricted settings, while taking into account the popularity of applying parameter efficient methods for fine-tuning, like LoRA or partial model freezing. In spirit, our work is similar to AutoML [He et al., 2021b], which aims to automatically find the optimal ML architecture. However, our method goes beyond and targets the data and the training method as well.

Parameter-efficient fine-tuning Modern language models have a lot of parameters, and fine-tuning them using consumer-grade hardware is difficult in general. Notwithstanding the recent turn to *in-context learning*, fine-tuning is still necessary for most applications. A range of techniques have been developed recently, to make fine-tuning more efficient [Houlsby et al., 2019, Li and Liang, 2021, Qi et al., 2022, He et al., 2021a]. These approaches are all applicable to embedding models. Muennighoff [2022] explored a simple parameter-efficient approach to repurpose GPT models into embedding models where only the bias tensors of the transformer model are updated [Zaken et al., 2022]. [Sun et al., 2024] developed a parameter-efficient method designed specifically for fine-tuning embedding models. However, a systematic study of what parameter-efficient methods are optimal under what scenarios has not been performed, which is the aim of our work.

3 Preliminaries

3.1 Scaling laws and compute-optimal models

The constraint optimisation problem that we aim to solve is the following: *given a fine-tuning corpus and a family of pre-trained decoder-only language models of different sizes, minimise the contrastive loss subject to a fixed computational budget.*

Following Hoffmann et al. [2022, Section 3.2], we perform a hyperparameter search at different computational budget levels to obtain IsoFLOP profiles of models. We then find the optimal models at each level, fit a loss-size scaling law, and extrapolate it to unobserved FLOP budgets. We only focus on contrastive fine-tuning since state-of-the-art text embedder models are all trained in such fashion [Gao et al., 2021c, Neelakantan et al., 2022, Wang et al., 2023].

We refer to the models with the lowest contrastive loss achievable with fixed computational costs as *compute-optimal models*. We focus on optimising three components specifically: model size, data quantity, and fine-tuning method. For each fine-tuning method, we find the optimal model configuration at each FLOP level and establish the relationship between the computational budget and optimal loss for that method. Then, for any unobserved computational budget, we use the scaling laws to predict the loss, and pick the method that gives the lowest loss and its corresponding configuration.

3.2 Extracting representations from transformers

We start with a standard decoder-only transformer model architecture, such as GPT architectures [Radford et al., 2019, Brown et al., 2020], or Pythia [Biderman et al., 2023]. Given a sequence of tokens $S = (s_1, \dots, s_n)$, we pass them through the transformer model and average the last layer representations as the embedding vector $\text{emb}(S) = \frac{1}{n} \sum_{i=1}^n h_i \in \mathbb{R}^m$, where h_i is the last layer hidden state of the token s_i .

Some works apply the last token pooling approach [Neelakantan et al., 2022] instead of using mean pooling described above. Here, however, we default to mean pooling as in [Wang et al., 2022, Li et al., 2023, Günther et al., 2023]. It is a more popular method overall, and moreover, we find it to yield better performance, which we demonstrate in Appendix E.3.

3.3 Contrastive fine-tuning

We use the contrastive loss as in [Neelakantan et al., 2022], which has been widely used in fine-tuning pre-trained language models [Wang et al., 2022, 2023, Günther et al., 2023].

For a batch of n examples, that is, n pairs of texts $(x_1, y_1), \dots, (x_n, y_n)$, we only treat examples $(x_i, y_i)_{i=1}^n$ as positive matches, and all the pairs (x_i, y_j) , where $i \neq j$ as negatives. Concretely, we calculate the batch contrastive loss in two steps. First, we calculate the logits:

$$\text{logits}[i, j] = \text{sim}(\text{emb}(x_i), \text{emb}(y_j)) \cdot \exp(\tau),$$

where `sim` is the cosine similarity function between two vectors, and $\tau \in \mathbb{R}$ is the temperature. The total loss is the sum of the cross entropy losses on both the row and the column directions:¹

```
labels = [0, 1, ..., n - 1]
l_r = cross_entropy(logits, labels)
l_c = cross_entropy(logits.T, labels)
loss = (l_r + l_c) / 2
```

Hard negatives in contrastive learning In our contrastive learning setting, we use in-batch examples as negatives. Some datasets additionally include tailored negative examples (*hard negatives*) for each positive datapoint, which can be incorporated into the contrastive loss to promote learning more precise representations. Moreover, there are works that focus on approaches for mining hard negatives which result in better training data in the context of specific downstream tasks [Xiong et al., 2021, Zhan et al., 2021].

However, recent works aiming at training powerful, general-purpose embedding models that do rely on datasets with hard negatives often arrive at embeddings by having two distinct training phases: the expensive “unsupervised” *phase one* involving vast data from the internet, and then a “supervised” *phase two*, often targeted towards a specific downstream task, where training data of lower quantity – but containing high-quality hard negatives – is used [Li et al., 2023, Wang et al., 2022, Xiao et al., 2023]. Since the first phase is more expensive, it will benefit more from scaling laws such as the ones we derive.

Moreover, the usefulness of hard negatives highly depends on their quality, which may vary wildly between datasets. Therefore, conclusions reached with hard negatives are more closely tied to the datasets the models are trained on. Since we develop scaling laws agnostic to the dataset, we abstain from using hard negatives in our experiments.

3.4 Fine-tuning methods

The most basic and straightforward method of fine-tuning is *full fine-tuning*, where *all* the weights of the model are updated under the contrastive objective.

We also study parameter-efficient fine-tuning (PEFT) methods, which are popular, especially in academic settings, because they have lower memory requirements than full fine-tuning. Since the PEFT methods update fewer parameters in the network than there are in total, the relationship between their computational cost, model size, and data quantity is different from full fine-tuning. Hence, we individually study the IsoFLOP profiles for four fine-tuning methods, namely, *full fine-tuning*, *block freezing*, *bias-only tuning*, and *Low-Rank Adaptation (LoRA)*. Each of the last three non-standard methods is briefly characterised below.

Block freezing In this approach, we *freeze* the LLM token-embedding layer as well as the first k transformer blocks so that their parameters stay fixed during fine-tuning. Only the latter part of the model is back-propagated through and updates its parameters. To process the same amount of data, this is more computationally efficient than full fine-tuning. By varying the number of frozen blocks, one trades-off between the computational efficiency and flexibility of the model.

Low-rank adaptation (LoRA) This method was introduced by Hu et al. [2022] to significantly reduce the number of trainable parameters of a neural network while maintaining high performance. LoRA introduces a small number of weights into the model and only trains those. It can be applied to a subset of dense layers of a model. It works by parametrising each dense layer W from a subset of linear layers of a model by two low-rank matrices A and B , and using $W + AB$ instead of W , while training only A and B .

Bias-only tuning In this method, only bias parameters are fine-tuned, and the rest of the model remains frozen. This approach was proposed by Zaken et al. [2022] as a parameter-efficient fine-tuning method for BERT encoders, and recently applied by Muennighoff [2022] to contrastively fine-tune GPT models.

3.5 Calculating computational cost

We denote the number of non-token-embedding parameters used in the forward pass to process each token as N_F , the number of non-token-embedding parameters in backward-propagation as N_B , the num-

¹We adopt the PyTorch notation here, where the cross entropy loss takes two arguments: the logits of the probability distribution, and the indices of the true categories.

ber of non-token-embedding parameters updated during backward-propagation as N_U , the total number of tokens processed as D , and the computational cost as C floating point operations (FLOP). Following the calculation by Kaplan et al. [2020], we derive the relationship between the variables to be

$$C = 2N_F D + 2N_B D + 2N_U D.$$

In case of full fine-tuning, every parameter is trainable, so $N_F = N_B = N_U$, and therefore

$$C = 6N_F D.$$

This is the same as the formula for the computational cost used by Kaplan et al. [2020] for pre-training.

4 Experiments

We first specify the relevant details of our experimental setup (Section 4.1). Next, we present the results of our experiments where we contrastively train a grid of models of different sizes, using different computational budgets, and apply different compute-optimal fine-tuning methods with varying hyperparameters (Section 4.2). Based on the collected data, we fit a mathematical formula describing the observed scaling laws (Section 4.3). Finally, we provide general observations and recommendations for efficient contrastive fine-tuning (Section 4.6). In these experiments² we investigate the compute-constrained setup, but we present limited results on the data-constrained setup in Appendix D.

4.1 Experimental setup

We use eight decoder-only models from the Pythia suite [Biderman et al., 2023], which have sizes 14M, 31M, 70M, 160M, 410M, 1B, 1.4B, and 2.8B of parameters. All the models have been pre-trained on the Pile dataset [Gao et al., 2021a] for 300 billion tokens. We consider six computational budgets equally spaced on the log-scale: 1.5e15, 6e15, 2.4e16, 9.6e16, 3.8e17, and 1.5e18 FLOP.

We fine-tune our models on the English partition of the BAAI BGE dataset [Xiao et al., 2023], which contains 200 million semantically related (*query*, *value*) pairs from various internet sources such as Wikipedia and Stack Exchange. Note that since the BAAI BGE dataset is massive, for all our experiments we fine-tune for less than one epoch on this corpus, which allows us to avoid the effect of diminishing returns on subsequent training epochs [Muennighoff et al., 2023a].

We use the AdamW optimiser [Loshchilov and Hutter, 2019] and a cosine learning rate scheduler during training. The learning rate first goes through a linear warm-up phase of 1/10 of the total steps to a peak that is 1/10 of the pre-training peak learning rate, that is, to learning rates between 1.2×10^{-5} and 6×10^{-5} . Then, it decays in a cosine schedule to 1/10 of the maximum at the end of training. We employ gradient checkpointing [Chen et al., 2016] and GradCache [Gao et al., 2021b] to limit the peak GPU RAM usage during training. We use the Transformers [Wolf et al., 2020] library for the training of the models and use the Accelerate [Gugger et al., 2022] library to facilitate multi-GPU training. The batch size in our main line of experiments is 1024 sequences and the context length is 75 tokens. We give the full list of training hyperparameters in Appendix E.

We also perform a less extensive grid of experiments to confirm that our findings are robust to change in hyperparameters. We find that our scaling law formula fitted to losses from models with batch size 1024 and context length of 75 tokens describes the loss for models trained with batch size 2048 and context length of 512 tokens very well. The detailed description and plots are shown in Appendix G.

In addition to controlling the final training contrastive loss achieved by the models, we also measure downstream performance by evaluating the models on a representative subset of the MTEB benchmark [Muennighoff et al., 2023b]. The benchmark defines eight categories of tasks in total (e.g., retrieval, semantic text similarity). We select one task from each category by determining which ones are the most correlated with the performance for the whole category (see Appendix B for details).

4.2 Experimental results for different methods

Full fine-tuning With full fine-tuning, the numbers of forward and backward parameters are the same, and the computational cost is straightforwardly $C = 6N_F D$. In Figure 2a, we present the final contrastive loss vs. the number of parameters in the network on a log-log scale.

²The data collected in our experiments and their interactive visualisations are available in this Colab notebook: <https://colab.research.google.com/drive/1EC2QdVrOIXhamLJI51CUywWJ33sApXX8>

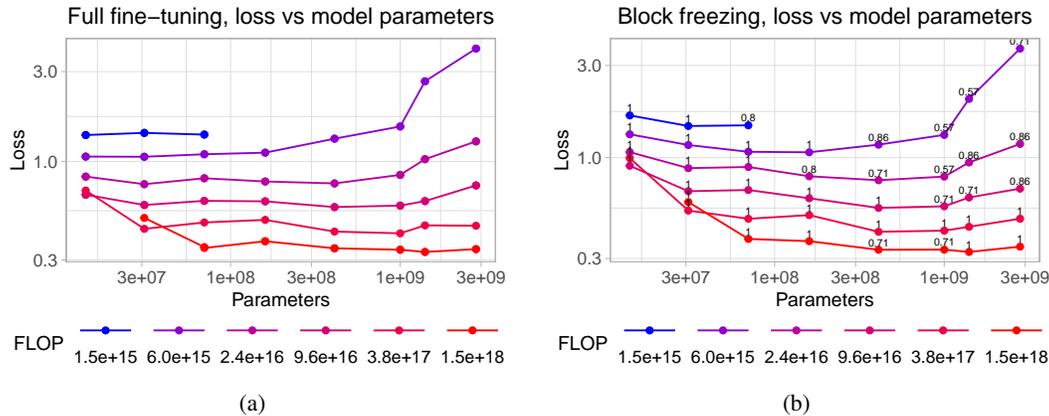


Figure 2: **(a)** IsoFLOP profiles for *full fine-tuning*. The horizontal axis is the number of parameters in the model, and the vertical axis is the achieved loss. Both axes use log-scale. The optimal model size tends to increase as the computational budget increases. **(b)** IsoFLOP profiles for *block freezing*. The axes are the same as for full fine-tuning. Each data point denotes the optimal choice with respect to the fraction of active blocks during training, which is noted above the points. The optimal model size tends to increase as the computational budget increases, while the optimal active block fraction tends to slightly decrease as the model size gets larger.

As expected, larger computational budgets consistently improve the loss. The IsoFLOP profiles resemble the classical ones from Hoffmann et al. [2022]. However, they tend to be flatter when the model size and computational budget are both small or both large. From the IsoFLOP profiles we can see what is the optimal model size that minimises the loss at each computational budget.

Block freezing Since the Pythia models we experimented with have 6, 12, 16, 24 or 32 blocks, we freeze $\frac{0}{6}, \frac{1}{6}, \dots, \frac{5}{6}$ of their blocks if their number is divisible by 6, and $\frac{0}{8}, \frac{1}{8}, \dots, \frac{7}{8}$ if it is divisible by 8. In this setup, unlike for full fine-tuning, the token-embedding parameters always remain frozen even when all the transformer blocks are active. We have $N_B = N_U < N_F$, and hence cost $C = 2N_FD + 4N_BD$.

In Figure 2b, we present IsoFLOP profiles for optimal loss at given model sizes. The trend of the profiles is quite similar to that of full fine-tuning presented in Figure 2a.

We plot the loss with varying model sizes, computational budgets, and fractions of frozen blocks in Figure 3. For smaller models ($\leq 160M$ parameters), a lower fraction of frozen blocks leads to a lower final loss across all computational budgets. However, for larger models, using more than 70% of transformer often only marginally improves the model. For instance, for the model with 1e9 parameters, it is visible that freezing 50 – 70% of the blocks gives the optimal results for all budgets except for 3.8e17 FLOP.

Bias-only tuning Here we simply update only the bias parameters and not the weights of the model. Figure 4a shows the final training loss across all the model sizes and computational budgets. Increasing the computational budget decreases the optimal loss that is achievable, although the absolute values of the losses are high compared to other fine-tuning methods (the lowest achievable loss is above 0.4 with bias-only tuning).

Low-rank adaptation We use the LoRA implementation from the PEFT library [Mangrulkar et al., 2022]. We attach the adapter to all the dense layers of the network, and vary the adapter rank from 8 up to 2048, since even lower ranks resulted in an unstable training.

We compare the loss obtained for each model size and computational budget for different LoRA rank values in Figure 5. In almost all the combinations of the model size and computational budget, the rank of 32 or 128 (occasionally 512) is optimal. Wang et al. [2023] also use LoRA in the context of contrastive learning, but they fix the rank of 16 for their experiments. Our findings suggest that this choice should be reconsidered as being potentially suboptimal.

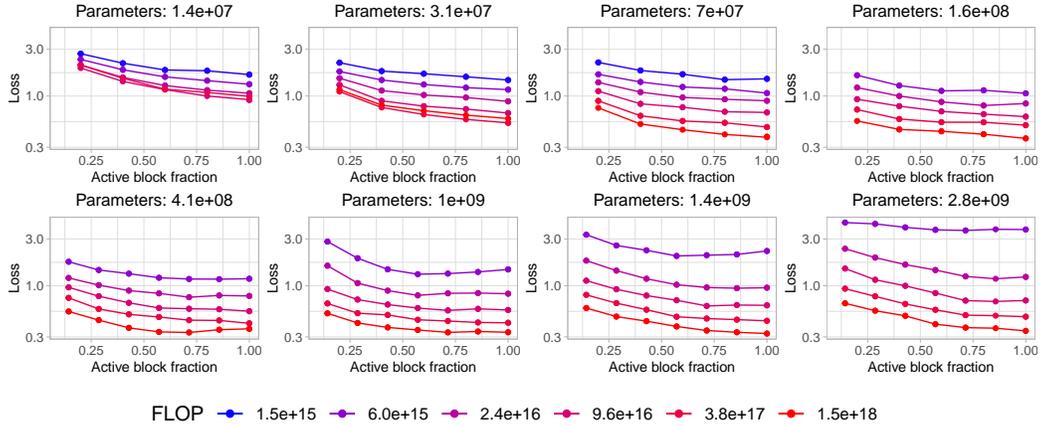


Figure 3: The effect of block freezing across all model sizes. Different colours signify different computational budgets. Unless the model is large and the computational budget small, it is always better to update all the (non-embedding) weights of the model.

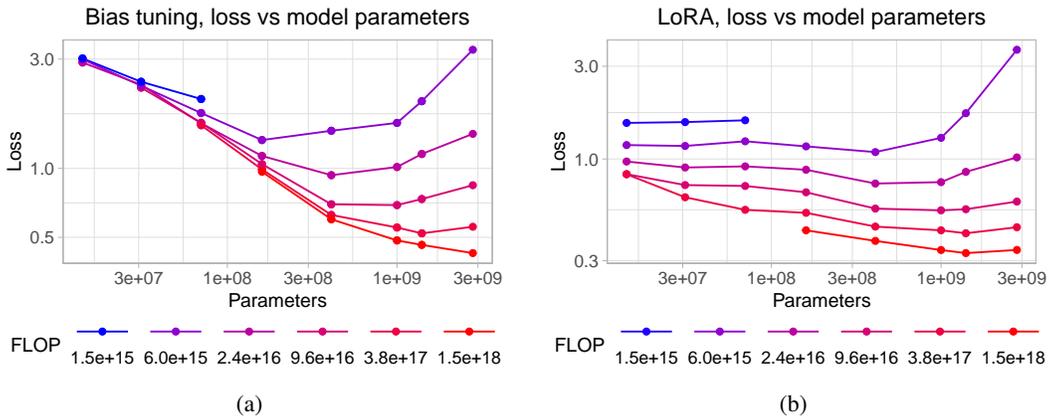


Figure 4: (a) IsoFLOP profiles for *bias-only tuning*. The horizontal axis is the number of parameters in the model, and the vertical axis is the achieved loss. Both axes use log-scale. The optimal model size increases as the computational budget increases, but the achievable loss is higher than for other fine-tuning methods. (b) IsoFLOP profiles for *LoRA* fine-tuning. The axes are the same as for bias-only tuning. Each data point denotes the optimal choice of the rank of LoRA matrices given the size of the model and the computational budget.

In Figure 4b, we present IsoFLOP profiles for LoRA, where each data point represents the best LoRA rank setting for a given model size and computational budget. We see that the relationship between the optimal loss and the LoRA rank is quite consistent across model sizes and computational budgets, meaning that the rank hyperparameter is not very sensitive with respect to the two variables. In practice, this means that setting the LoRA rank to 32 or 128 are likely to be good default choices.

Lookup table of optimal model sizes for different methods In Figure 6a (for full fine-tuning), Figure 6b (for LoRA), and Figure 13 (for freezing and bias tuning, in Appendix A), we visualise the relationship between the optimal model size and the computational budget. This makes looking up the optimal configuration for a given method at a given budget very convenient.

Downstream performance In Figures 8, 9, 10, and 11 (in Appendix A) we present the IsoFLOP plots of the downstream performance for all the four fine-tuning methods on the subset of the MTEB benchmark (described in Appendix B). Based on the performance results, we also present the optimal size choices per budget in Figure 15 in Appendix A (which is a performance analogue of Figure 13), and the optimal performance plot in Figure 12 (a performance analogue of Figure 1). These are mostly consistent with the loss results, albeit more noisy. In Appendix C, we demonstrate the significant correlation between the final training loss and the downstream task performance.

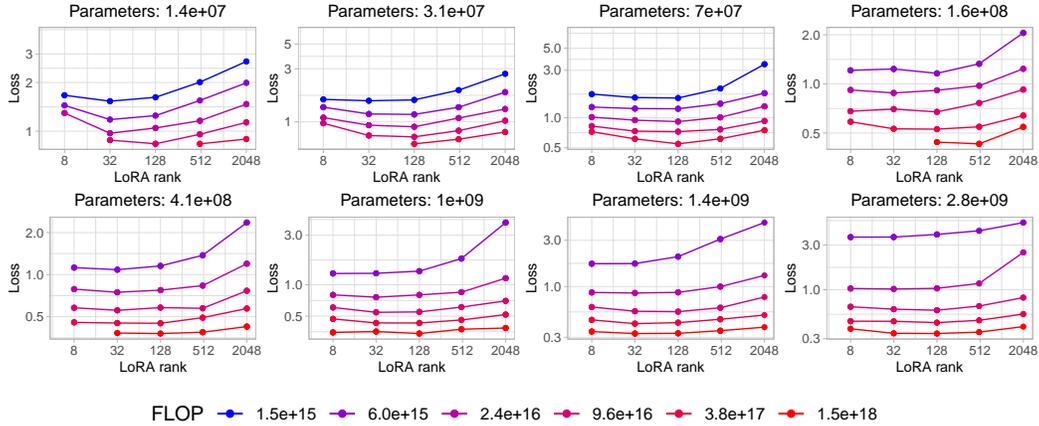


Figure 5: The effect of different LoRA ranks across all model sizes. Different colours signify different computational budgets. The inflected curves indicate that it is less beneficial to use a rank from either extremes of the spectrum (8 or 2048). The detrimental effect of the high rank of 2048 is stronger for lower computational budgets. Ranks of 32 and 128 result in the lowest loss overall.

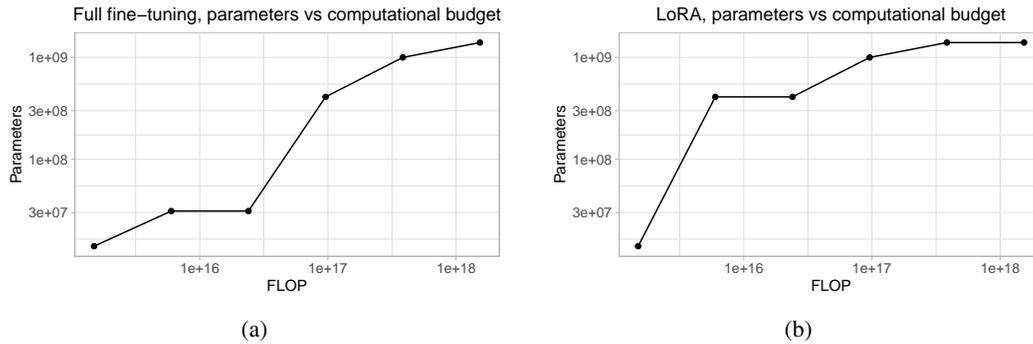


Figure 6: Optimal model size vs. computational budget for (a) full fine-tuning, and (b) LoRA.

4.3 Scaling laws for embeddings

We aim to model the behaviour of the loss L with an analytical formula as a function of the number of parameters N and the number of training tokens D . Following Hoffmann et al. [2022], we start with the following formula:

$$L(N, D) = C + \frac{A}{N^\alpha} + \frac{B}{D^\beta},$$

where $A, B, C, \alpha, \beta \in \mathbb{R}_+$ are real-valued parameters to be fitted. We found that this formula reasonably describes the loss behaviour for each of our methods. However, the fraction of *trainable parameters* is an important factor that is not considered in the model above. Following the modelling of Frantar et al. [2023], we propose a modified formula:

$$L(S, N, D) = C + \frac{a_d \log(D) + b_d}{N^\alpha} + \frac{a_s(1 - S)^{b_s} + c_s}{D^\beta},$$

where S is an additional variable representing the trainable fraction of parameters, and $a_d, b_d, a_s, b_s, c_s \in \mathbb{R}$ are additional coefficients to be fitted. We check that this new formula fits the data better than the original one, and conjecture that it can be used for larger models and larger computational budgets. More details about the derivation are presented in Appendix F.

4.4 Compute-optimal frontier and recipe

In Figure 1, we plotted the best achievable losses against the computational budgets for the four fine-tuning methods considered. We then fitted a linear trend for each method on a log-log scale (equivalent to a power law relationship on a normal scale), and highlighted the lowest loss achievable at given

Algorithm 1 Recipe for compute-optimal embedding model

Input: compute budget C .

Output: fine-tuning method, model size, data quantity, and (optionally) the method's hyperparams.

if $B < 9.06e16$ FLOP **then**

 Use full fine-tuning.

 Go to Figure 6a to find the optimal model parameters N given budget C .

 Calculate the data quantity D .

return Full fine-tuning, N , D , ().

else

 Use LoRA.

 Go to Figure 6b to find the optimal model parameters N given budget C .

 Go to Figure 5 to find the LoRA rank R according to C and N .

return LoRA, N , D , (R).

budgets using the best corresponding fine-tuning method. We call the contour of this function that predicts the optimal loss across methods the compute-optimal frontier.

The equation describing the full fine-tuning and the LoRA linear fits are

$$\ln(\text{loss}) = -0.21 \cdot \ln(\text{budget}) + 8.39, \text{ and } \ln(\text{loss}) = -0.22 \cdot \ln(\text{budget}) + 8.93,$$

respectively. The equations suggest that at a budget lower than 9.06e16 FLOP, the optimal model is achieved with full fine-tuning, and at a higher budget, the optimal model is achieved with LoRA.

With this, one becomes fully equipped to deduce the optimal recipe for text embedding model training for a given budget, by carrying out the procedure detailed in Algorithm 1. Note that although the block freezing approach is not predicted to be optimal at any budget, its performance is quite close to optimal. It also has the advantage of being easy to implement and has lower memory requirements, hence might be a good method to choose in practical use cases for larger budgets. Its optimal model size and fraction of active blocks hyperparameter can be worked out in the same way as for the other methods, with the help of Figure 3 and Figure 13 (left pane) in Appendix A.

We briefly discuss the case of data-optimal frontier in Appendix D. We do not emphasise this setup due to it being less standard in language modelling research.

4.5 Generalisation

To assess the generality of our findings, we investigate whether our analytical formula for predicting loss—fitted to datapoints from the Pythia suite (from Section 4.3) can also describe the behaviour of models from a different family. We conduct experiments on two language models, namely Gemma 2B [Mesnard et al., 2024] and Gemma 2 2B Team et al. [2024]. We trained both models with the computational budgets of 1.5e15, 6e15, 2.4e16, 9.6e16, 3.8e17, and 1.5e18 FLOP.

Figure 7 presents results for Gemma 2B with both full fine-tuning and LoRA fine-tuning, while results for Gemma 2 2B are shown in Appendix A. In both cases, the formula provides a close approximation of the loss behavior. Although this study is limited to a single model size, it offers an initial indication that our findings may generalize to new models.

4.6 Takeaways

We compile the takeaways from our extensive experiments by analysing each fine-tuning method individually and contrasting them afterwards.

- Bias-only tuning is not a good fine-tuning strategy for embedding model training, as it is consistently worse than other strategies under IsoFLOP comparisons.
- LoRA and block freezing are both effective approaches that improve the performance of the trained embedding models at higher budgets. For LoRA, the rank hyperparameter is not very sensitive to model size or computational budgets, with an optimum at around 128.
- Finding the optimal recipe for fine-tuning embedding models requires a non-trivial effort, and the resulting optimal configurations are sometimes counter-intuitive (e.g., not using the largest model size possible). Carefully tuning hyperparameters is crucial and should be practised systematically.

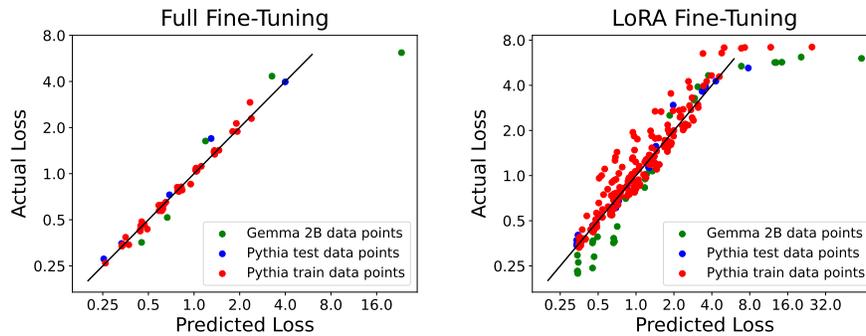


Figure 7: Actual loss vs. predicted loss for full fine-tuning and LoRA finetuning. Red points are the models from runs on Pythia suite of size smaller than 2.8B, blue points are from runs on Pythia 2.8B and green points are from runs on Gemma 2B. The formula is fitted using the red points.

5 Limitations and future work

Other families of models In this paper, we experimented with the Pythia [Biderman et al., 2023] suite of pre-trained models, which are of different sizes but have all been trained on 300B tokens. This means that they have been over-trained with respect to the Chinchilla scaling law [Hoffmann et al., 2022] to different degrees. The relationship between the optimal model size and the computational budget might differ for a suite of pre-trained models with different over-training ratios. We note that prior efforts [Dettmers et al., 2024] have found that fine-tuning conclusions reached with Pythia do generalise to other model families such as OPT [Zhang et al., 2022], LLaMA [Touvron et al., 2023], and BLOOM [Scao et al., 2022], in the quantised setting. Indeed, in Section 4.5, we find that a model from a different family (Gemma [Mesnard et al., 2024]) closely fits our loss predictions.

Averaging Due to constraints on our computational resources, we ran each experiment exactly once. Averaging over multiple random seeds for the experiments would reduce the variance involved in each, and potentially result in less noisy conclusions. This is a straightforward extension that we leave to carry out in future work, should we have more resources. We also plan to evaluate the trained models on the entirety of the MTEB benchmark [Muennighoff et al., 2023b], instead of using the losses from the last steps of training, to better reflect the performances of embedding models.

Other forms of embedding readout We only experimented with mean pooling as a way of extracting embeddings from the transformer models. Other readout methods could be used, like the max pooling or extracting the hidden state corresponding to the end-of-sequence token, which might result in different scaling laws and optimal frontiers.

Inference cost analysis In this work, we only focus on the training cost and refrain from taking inference costs into account. However, in certain practical scenarios, the latter may also be relevant.

6 Conclusions

In this paper, we systematically investigated the influence of pre-trained model size, fine-tuning data quantity, and fine-tuning method in the final performance of embedding models repurposed from language models. As a result, we devised an algorithm that takes in the computational budget and returns the optimal configuration for the embedding model to train. This enables researchers who wish to adapt language models to embed their own data, especially those with limited computational resources, to obtain optimal text embedding models with greater efficiency in time and resources.

Acknowledgements

The authors would like to thank Denis Kuznedelev, Soroush Tabesh and Konrad Staniszewski for helpful discussions and feedback. AQJ is supported by a Peterhouse graduate studentship. PM was supported by National Science Center Poland under the grant agreement 2019/35/O/ST6/03464. We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Centers: ACK Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2023/016717.

References

- R. Anil, A. M. Dai, O. Firat, M. Johnson, D. Lepikhin, A. Passos, S. Shakeri, E. Taropa, P. Bailey, Z. Chen, E. Chu, J. H. Clark, L. E. Shafey, Y. Huang, K. Meier-Hellstern, G. Mishra, E. Moreira, M. Omernick, K. Robinson, S. Ruder, Y. Tay, K. Xiao, Y. Xu, Y. Zhang, G. H. Ábrego, J. Ahn, J. Austin, P. Barham, J. A. Botha, J. Bradbury, S. Brahma, K. Brooks, M. Catasta, Y. Cheng, C. Cherry, C. A. Choquette-Choo, A. Chowdhery, C. Crepy, S. Dave, M. Dehghani, S. Dev, J. Devlin, M. Díaz, N. Du, E. Dyer, V. Feinberg, F. Feng, V. Fienber, M. Freitag, X. Garcia, S. Gehrmann, L. Gonzalez, and et al. Palm 2 technical report. *CoRR*, abs/2305.10403, 2023. doi: 10.48550/ARXIV.2305.10403. URL <https://doi.org/10.48550/arXiv.2305.10403>.
- D. Biderman, J. G. Ortiz, J. Portes, M. Paul, P. Greengard, C. Jennings, D. King, S. Havens, V. Chiley, J. Frankle, C. Blakeney, and J. P. Cunningham. LoRA learns less and forgets less, 2024.
- S. Biderman, H. Schoelkopf, Q. G. Anthony, H. Bradley, K. O’Brien, E. Hallahan, M. A. Khan, S. Purohit, U. S. Prashanth, E. Raff, A. Skowron, L. Sutawika, and O. van der Wal. Pythia: A suite for analyzing large language models across training and scaling. In A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, editors, *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 2397–2430. PMLR, 2023. URL <https://proceedings.mlr.press/v202/biderman23a.html>.
- T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. Language models are few-shot learners. *CoRR*, abs/2005.14165, 2020. URL <https://arxiv.org/abs/2005.14165>.
- T. Chen, B. Xu, C. Zhang, and C. Guestrin. Training deep nets with sublinear memory cost. *CoRR*, abs/1604.06174, 2016. URL <http://arxiv.org/abs/1604.06174>.
- T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer. QLoRA: Efficient finetuning of quantized LLMs. *Advances in Neural Information Processing Systems*, 36, 2024.
- J. Devlin, M. Chang, K. Lee, and K. Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In J. Burstein, C. Doran, and T. Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics, 2019. doi: 10.18653/V1/N19-1423. URL <https://doi.org/10.18653/v1/n19-1423>.
- Y. Fang, J. Zhan, Q. Ai, J. Mao, W. Su, J. Chen, and Y. Liu. Scaling laws for dense retrieval, 2024. URL <https://arxiv.org/abs/2403.18684>.
- E. Frantar, C. Riquelme, N. Houlsby, D. Alistarh, and U. Evci. Scaling laws for sparsely-connected foundation models. *CoRR*, abs/2309.08520, 2023. doi: 10.48550/ARXIV.2309.08520. URL <https://doi.org/10.48550/arXiv.2309.08520>.
- L. Gao, S. Biderman, S. Black, L. Golding, T. Hoppe, C. Foster, J. Phang, H. He, A. Thite, N. Nabeshima, S. Presser, and C. Leahy. The pile: An 800GB dataset of diverse text for language modeling. *CoRR*, abs/2101.00027, 2021a. URL <https://arxiv.org/abs/2101.00027>.
- L. Gao, Y. Zhang, J. Han, and J. Callan. Scaling deep contrastive learning batch size under memory limited setup. In A. Rogers, I. Calixto, I. Vulic, N. Saphra, N. Kassner, O. Camburu, T. Bansal, and V. Shwartz, editors, *Proceedings of the 6th Workshop on Representation Learning for NLP, RepL4NLP@ACL-IJCNLP 2021, Online, August 6, 2021*, pages 316–321. Association for Computational Linguistics, 2021b. doi: 10.18653/V1/2021.REPL4NLP-1.31. URL <https://doi.org/10.18653/v1/2021.rep14nlp-1.31>.
- T. Gao, X. Yao, and D. Chen. SimCSE: Simple contrastive learning of sentence embeddings. In M. Moens, X. Huang, L. Specia, and S. W. Yih, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 6894–6910. Association for Computational

- Linguistics, 2021c. doi: 10.18653/V1/2021.EMNLP-MAIN.552. URL <https://doi.org/10.18653/v1/2021.emnlp-main.552>.
- S. Gugger, L. Debut, T. Wolf, P. Schmid, Z. Mueller, S. Mangrulkar, M. Sun, and B. Bossan. Accelerate: Training and inference at scale made simple, efficient and adaptable. <https://github.com/huggingface/accelerate>, 2022.
- M. Günther, L. Milliken, J. Geuter, G. Mastrapas, B. Wang, and H. Xiao. Jina embeddings: A novel set of high-performance sentence embedding models. *CoRR*, abs/2307.11224, 2023. doi: 10.48550/ARXIV.2307.11224. URL <https://doi.org/10.48550/arXiv.2307.11224>.
- J. He, C. Zhou, X. Ma, T. Berg-Kirkpatrick, and G. Neubig. Towards a unified view of parameter-efficient transfer learning. *CoRR*, abs/2110.04366, 2021a. URL <https://arxiv.org/abs/2110.04366>.
- X. He, K. Zhao, and X. Chu. AutoML: A survey of the state-of-the-art. *Knowledge-Based Systems*, 212:106622, 2021b.
- J. Hoffmann, S. Borgeaud, A. Mensch, E. Buchatskaya, T. Cai, E. Rutherford, D. de Las Casas, L. A. Hendricks, J. Welbl, A. Clark, T. Hennigan, E. Noland, K. Millican, G. van den Driessche, B. Damoc, A. Guy, S. Osindero, K. Simonyan, E. Elsen, J. W. Rae, O. Vinyals, and L. Sifre. Training compute-optimal large language models. *CoRR*, abs/2203.15556, 2022. doi: 10.48550/ARXIV.2203.15556. URL <https://doi.org/10.48550/arXiv.2203.15556>.
- N. Hounsby, A. Giurgiu, S. Jastrzebski, B. Morrone, Q. de Laroussilhe, A. Gesmundo, M. Attariyan, and S. Gelly. Parameter-efficient transfer learning for NLP. *CoRR*, abs/1902.00751, 2019. URL <http://arxiv.org/abs/1902.00751>.
- E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen. LoRA: Low-rank adaptation of large language models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. URL <https://openreview.net/forum?id=nZeVKeeFYf9>.
- G. Izacard, M. Caron, L. Hosseini, S. Riedel, P. Bojanowski, A. Joulin, and E. Grave. Unsupervised dense information retrieval with contrastive learning. *Trans. Mach. Learn. Res.*, 2022, 2022. URL <https://openreview.net/forum?id=jKN1pXi7b0>.
- A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. de Las Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, L. R. Lavaud, M. Lachaux, P. Stock, T. L. Scao, T. Lavril, T. Wang, T. Lacroix, and W. E. Sayed. Mistral 7B. *CoRR*, abs/2310.06825, 2023. doi: 10.48550/ARXIV.2310.06825. URL <https://doi.org/10.48550/arXiv.2310.06825>.
- J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei. Scaling laws for neural language models. *CoRR*, abs/2001.08361, 2020. URL <https://arxiv.org/abs/2001.08361>.
- B. Lester, R. Al-Rfou, and N. Constant. The power of scale for parameter-efficient prompt tuning. In M. Moens, X. Huang, L. Specia, and S. W. Yih, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 3045–3059. Association for Computational Linguistics, 2021. doi: 10.18653/V1/2021.EMNLP-MAIN.243. URL <https://doi.org/10.18653/v1/2021.emnlp-main.243>.
- X. L. Li and P. Liang. Prefix-tuning: Optimizing continuous prompts for generation. *CoRR*, abs/2101.00190, 2021. URL <https://arxiv.org/abs/2101.00190>.
- Z. Li, X. Zhang, Y. Zhang, D. Long, P. Xie, and M. Zhang. Towards general text embeddings with multi-stage contrastive learning. *CoRR*, abs/2308.03281, 2023. doi: 10.48550/ARXIV.2308.03281. URL <https://doi.org/10.48550/arXiv.2308.03281>.
- I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=Bkg6RiCqY7>.

- S. Mangrulkar, S. Gugger, L. Debut, Y. Belkada, S. Paul, and B. Bossan. PEFT: State-of-the-art parameter-efficient fine-tuning methods. <https://github.com/huggingface/peft>, 2022.
- R. Meng, Y. Liu, S. R. Joty, C. Xiong, Y. Zhou, and S. Yavuz. SFR-Embedding-Mistral: Enhance text retrieval with transfer learning. Salesforce AI Research Blog, 2024. URL <https://blog.salesforceairesearch.com/sfr-embedded-mistral/>.
- T. Mesnard, C. Hardin, R. Dadashi, S. Bhupatiraju, S. Pathak, L. Sifre, M. Rivière, M. S. Kale, J. Love, P. Tafti, L. Hussenot, A. Chowdhery, A. Roberts, A. Barua, A. Botev, A. Castro-Ros, A. Slone, A. Héliou, A. Tacchetti, A. Bulanova, A. Paterson, B. Tsai, B. Shahriari, C. L. Lan, C. A. Choquette-Choo, C. Crepy, D. Cer, D. Ippolito, D. Reid, E. Buchatskaya, E. Ni, E. Noland, G. Yan, G. Tucker, G. Muraru, G. Rozhdestvenskiy, H. Michalewski, I. Tenney, I. Grishchenko, J. Austin, J. Keeling, J. Labanowski, J. Lespiau, J. Stanway, J. Brennan, J. Chen, J. Ferret, J. Chiu, and et al. Gemma: Open models based on gemini research and technology. *CoRR*, abs/2403.08295, 2024. doi: 10.48550/ARXIV.2403.08295. URL <https://doi.org/10.48550/arXiv.2403.08295>.
- T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. In Y. Bengio and Y. LeCun, editors, *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*, 2013. URL <http://arxiv.org/abs/1301.3781>.
- N. Muennighoff. SGPT: GPT sentence embeddings for semantic search. *CoRR*, abs/2202.08904, 2022. URL <https://arxiv.org/abs/2202.08904>.
- N. Muennighoff, A. M. Rush, B. Barak, T. L. Scao, A. Piktus, N. Tazi, S. Pyysalo, T. Wolf, and C. Raffel. Scaling data-constrained language models. *CoRR*, abs/2305.16264, 2023a. doi: 10.48550/ARXIV.2305.16264. URL <https://doi.org/10.48550/arXiv.2305.16264>.
- N. Muennighoff, N. Tazi, L. Magne, and N. Reimers. MTEB: Massive text embedding benchmark. In A. Vlachos and I. Augenstein, editors, *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2023, Dubrovnik, Croatia, May 2-6, 2023*, pages 2006–2029. Association for Computational Linguistics, 2023b. doi: 10.18653/V1/2023.EACL-MAIN.148. URL <https://doi.org/10.18653/v1/2023.eacl-main.148>.
- A. Neelakantan, T. Xu, R. Puri, A. Radford, J. M. Han, J. Tworek, Q. Yuan, N. Tezak, J. W. Kim, C. Hallacy, J. Heidecke, P. Shyam, B. Power, T. E. Nekoul, G. Sastry, G. Krueger, D. Schnurr, F. P. Such, K. Hsu, M. Thompson, T. Khan, T. Sherbakov, J. Jang, P. Welinder, and L. Weng. Text and code embeddings by contrastive pre-training. *CoRR*, abs/2201.10005, 2022. URL <https://arxiv.org/abs/2201.10005>.
- W. Qi, Y. Ruan, Y. Zuo, and T. Li. Parameter-efficient tuning on layer normalization for pre-trained language models. *CoRR*, abs/2211.08682, 2022. doi: 10.48550/ARXIV.2211.08682. URL <https://doi.org/10.48550/arXiv.2211.08682>.
- A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. Language models are unsupervised multitask learners. 2019.
- T. L. Scao, A. Fan, C. Akiki, E. Pavlick, S. Ilic, D. Hesslow, R. Castagné, A. S. Luccioni, F. Yvon, M. Gallé, J. Tow, A. M. Rush, S. Biderman, A. Webson, P. S. Ammanamanchi, T. Wang, B. Sagot, N. Muennighoff, A. V. del Moral, O. Ruwase, R. Bawden, S. Bekman, A. McMillan-Major, I. Beltagy, H. Nguyen, L. Saulnier, S. Tan, P. O. Suarez, V. Sanh, H. Laurençon, Y. Jernite, J. Launay, M. Mitchell, C. Raffel, A. Gokaslan, A. Simhi, A. Soroa, A. F. Aji, A. Alfassy, A. Rogers, A. K. Nitzav, C. Xu, C. Mou, C. Emezue, C. Klamm, C. Leong, D. van Strien, D. I. Adelani, and et al. BLOOM: A 176b-parameter open-access multilingual language model. *CoRR*, abs/2211.05100, 2022. doi: 10.48550/ARXIV.2211.05100. URL <https://doi.org/10.48550/arXiv.2211.05100>.
- S. Sun, H. Zhang, Z. Liu, J. Bao, and D. Song. Llm-oriented retrieval tuner, 2024. URL <https://arxiv.org/abs/2403.01999>.
- G. Team, M. Riviere, S. Pathak, P. G. Sessa, C. Hardin, S. Bhupatiraju, L. Hussenot, T. Mesnard, B. Shahriari, A. Ramé, J. Ferret, P. Liu, P. Tafti, A. Friesen, M. Casbon, S. Ramos, R. Kumar, C. L. Lan, S. Jerome, A. Tsitsulin, N. Vieillard, P. Stanczyk, S. Girgin, N. Momchev, M. Hoffman,

- S. Thakoor, J.-B. Grill, B. Neyshabur, O. Bachem, A. Walton, A. Severyn, A. Parrish, A. Ahmad, A. Hutchison, A. Abdagic, A. Carl, A. Shen, A. Brock, A. Coenen, A. Laforge, A. Paterson, B. Bastian, B. Piot, B. Wu, B. Royal, C. Chen, C. Kumar, C. Perry, C. Welty, C. A. Choquette-Choo, D. Sinopalnikov, D. Weinberger, D. Vijaykumar, D. Rogozińska, D. Herbison, E. Bandy, E. Wang, E. Noland, E. Moreira, E. Senter, E. Eltyshev, F. Visin, G. Rasskin, G. Wei, G. Cameron, G. Martins, H. Hashemi, H. Klimczak-Plucińska, H. Batra, H. Dhand, I. Nardini, J. Mein, J. Zhou, J. Svensson, J. Stanway, J. Chan, J. P. Zhou, J. Carrasqueira, J. Iljazi, J. Becker, J. Fernandez, J. van Amersfoort, J. Gordon, J. Lipschultz, J. Newlan, J. yeong Ji, K. Mohamed, K. Badola, K. Black, K. Millican, K. McDonell, K. Nguyen, K. Sodhia, K. Greene, L. L. Sjoesund, L. Usui, L. Sifre, L. Heuermann, L. Lago, L. McNealus, L. B. Soares, L. Kilpatrick, L. Dixon, L. Martins, M. Reid, M. Singh, M. Iverson, M. Görner, M. Velloso, M. Wirth, M. Davidow, M. Miller, M. Rahtz, M. Watson, M. Risdal, M. Kazemi, M. Moynihan, M. Zhang, M. Kahng, M. Park, M. Rahman, M. Khatwani, N. Dao, N. Bardoliwalla, N. Devanathan, N. Dumai, N. Chauhan, O. Wahltinez, P. Botarda, P. Barnes, P. Barham, P. Michel, P. Jin, P. Georgiev, P. Culliton, P. Kuppala, R. Comanescu, R. Merhej, R. Jana, R. A. Rokni, R. Agarwal, R. Mullins, S. Saadat, S. M. Carthy, S. Cogan, S. Perrin, S. M. R. Arnold, S. Krause, S. Dai, S. Garg, S. Sheth, S. Ronstrom, S. Chan, T. Jordan, T. Yu, T. Eccles, T. Hennigan, T. Kocisky, T. Doshi, V. Jain, V. Yadav, V. Meshram, V. Dharmadhikari, W. Barkley, W. Wei, W. Ye, W. Han, W. Kwon, X. Xu, Z. Shen, Z. Gong, Z. Wei, V. Cotruta, P. Kirk, A. Rao, M. Giang, L. Peran, T. Warkentin, E. Collins, J. Barral, Z. Ghahramani, R. Hadsell, D. Sculley, J. Banks, A. Dragan, S. Petrov, O. Vinyals, J. Dean, D. Hassabis, K. Kavukcuoglu, C. Farabet, E. Buchatskaya, S. Borgeaud, N. Fiedel, A. Joulin, K. Kenealy, R. Dadashi, and A. Andreev. Gemma 2: Improving open language models at a practical size, 2024. URL <https://arxiv.org/abs/2408.00118>.
- N. Thakur, N. Reimers, A. Rücklé, A. Srivastava, and I. Gurevych. BEIR: A heterogeneous benchmark for zero-shot evaluation of information retrieval models. In J. Vanschoren and S. Yeung, editors, *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*, 2021. URL <https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/65b9eeae6e1cc6bb9f0cd2a47751a186f-Abstract-round2.html>.
- H. Touvron, T. Lavril, G. Izacard, X. Martinet, M. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, and G. Lample. Llama: Open and efficient foundation language models. *CoRR*, abs/2302.13971, 2023. doi: 10.48550/ARXIV.2302.13971. URL <https://doi.org/10.48550/arXiv.2302.13971>.
- L. Wang, N. Yang, X. Huang, B. Jiao, L. Yang, D. Jiang, R. Majumder, and F. Wei. Text embeddings by weakly-supervised contrastive pre-training. *CoRR*, abs/2212.03533, 2022. doi: 10.48550/ARXIV.2212.03533. URL <https://doi.org/10.48550/arXiv.2212.03533>.
- L. Wang, N. Yang, X. Huang, L. Yang, R. Majumder, and F. Wei. Improving text embeddings with large language models. *CoRR*, abs/2401.00368, 2023. URL <https://arxiv.org/abs/2401.00368>.
- T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online, Oct. 2020. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.
- S. Xiao, Z. Liu, P. Zhang, and N. Muennighof. C-pack: Packaged resources to advance general chinese embedding. *CoRR*, abs/2309.07597, 2023. doi: 10.48550/ARXIV.2309.07597. URL <https://doi.org/10.48550/arXiv.2309.07597>.
- L. Xiong, C. Xiong, Y. Li, K. Tang, J. Liu, P. N. Bennett, J. Ahmed, and A. Overwijk. Approximate nearest neighbor negative contrastive learning for dense text retrieval. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=zeFrfgYz1n>.
- E. B. Zaken, Y. Goldberg, and S. Ravfogel. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. In S. Muresan, P. Nakov, and A. Villavicencio,

editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, ACL 2022, Dublin, Ireland, May 22-27, 2022, pages 1–9. Association for Computational Linguistics, 2022. doi: 10.18653/V1/2022.ACL-SHORT.1. URL <https://doi.org/10.18653/v1/2022.acl-short.1>.

- J. Zhan, J. Mao, Y. Liu, J. Guo, M. Zhang, and S. Ma. Optimizing dense retrieval model training with hard negatives. In F. Diaz, C. Shah, T. Suel, P. Castells, R. Jones, and T. Sakai, editors, *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 1503–1512. ACM, 2021. doi: 10.1145/3404835.3462880. URL <https://doi.org/10.1145/3404835.3462880>.
- B. Zhang, Z. Liu, C. Cherry, and O. Firat. When scaling meets LLM finetuning: The effect of data, model and finetuning method. *CoRR*, abs/2402.17193, 2024. doi: 10.48550/ARXIV.2402.17193. URL <https://doi.org/10.48550/arXiv.2402.17193>.
- S. Zhang, S. Roller, N. Goyal, M. Artetxe, M. Chen, S. Chen, C. Dewan, M. Diab, X. Li, X. V. Lin, et al. OPT: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.

A Additional results

In Figure 8 (for full fine-tuning), Figure 9 (for block freezing), Figure 10 (for LoRA), and Figure 11 (for bias tuning), we present the IsoFLOP profiles where downstream task performance is shown. The downstream performance is measured on a subset of MTEB tasks described in Appendix B.

Using the location of optimal loss given the computational budget for freezing and bias tuning methods – as illustrated in Figures 2b and 4a – we project optimal model sizes for each computational budget in Figure 13. Similar figures for full fine-tuning and LoRA are presented in the main text in Figure 6a and Figure 6b, respectively.

Using the location of *optimal downstream task performance* given the computational budget for all the four fine-tuning methods – as illustrated in Figures 8, 9, 10, and 11 – we project optimal model sizes for each computational budget in Figure 15.

In Figure 14 we present how our formula fitted to models from the Pythia suite describes the loss behaviour of the Gemma 2 model.

B Downstream task evaluation details

For each category present in the MTEB benchmark [Muennighoff et al., 2023b], we take the data from the models evaluated in the original paper and select the task that has the highest correlation with the category average. This results in 8 tasks from the MTEB benchmark used for the downstream task evaluation:

1. EmotionClassification (classification),
2. TwentyNewsgroupsClustering (clustering),
3. SprintDuplicateQuestions (pair classification),
4. AskUbuntuDupQuestions (re-ranking),
5. SciFact (retrieval),
6. STS15 (semantic text similarity)
7. SummEval (summarization),
8. Tatoeba (bitext mining).

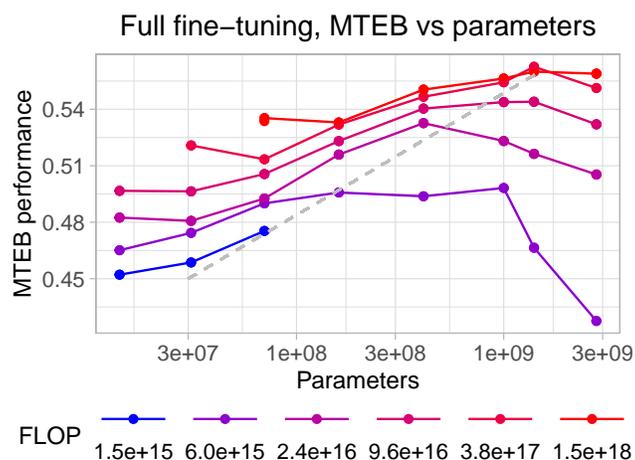
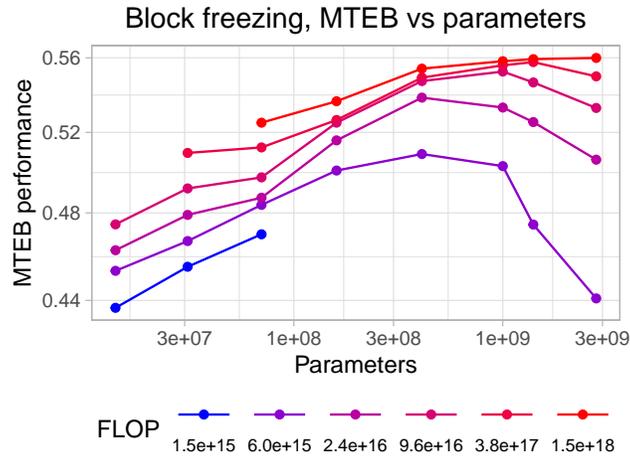


Figure 8: IsoFLOP profiles for *full fine-tuning*, where performance on MTEB tasks is observed instead of the loss. Both axes use log-scale. The performance remains approximately inversely correlated with the loss. However, interestingly, there are no corresponding plateaus on the IsoFLOP profiles.



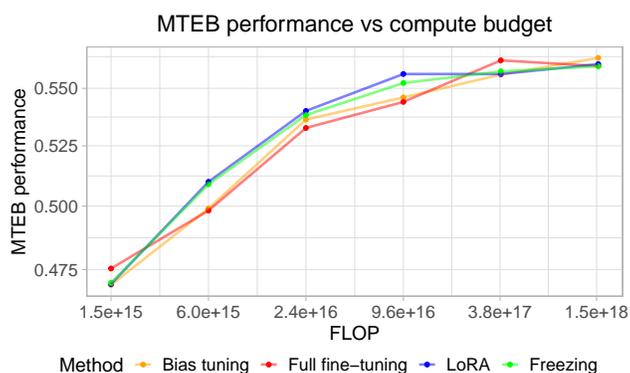


Figure 12: The optimal downstream task performance achieved using all four different fine-tuning methods at given budgets. This is a performance analogue of Figure 1 (which shows final training loss instead).

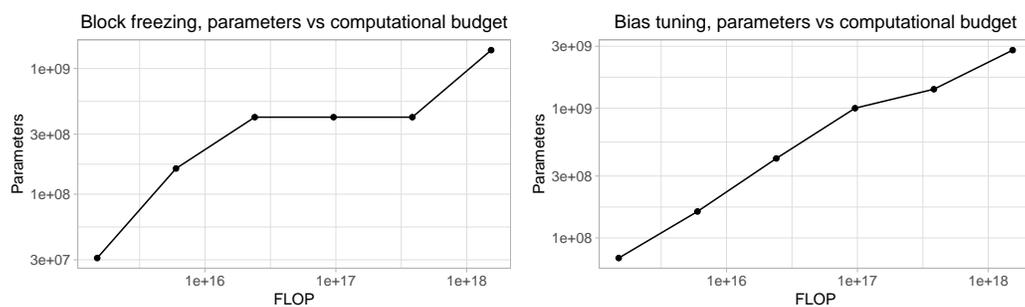


Figure 13: Optimal choices of model size given the computational budget, with respect to the *final training loss*, for block freezing and bias tuning.

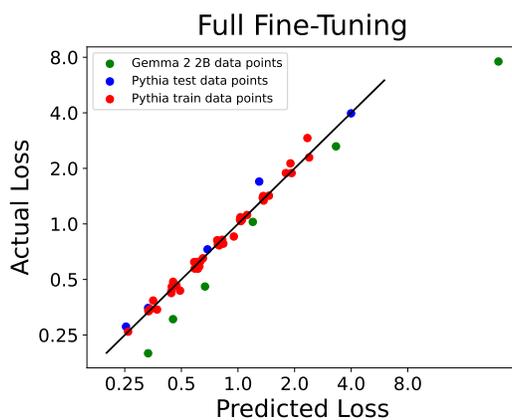


Figure 14: Actual loss vs. predicted loss for full fine-tuning and LoRA finetuning. Red points are the models from runs on Pythia suite of size smaller than 2.8B, blue points are from runs on Pythia 2.8B and green points are from runs on Gemma 2 2B. The formula is fitted using the red points.

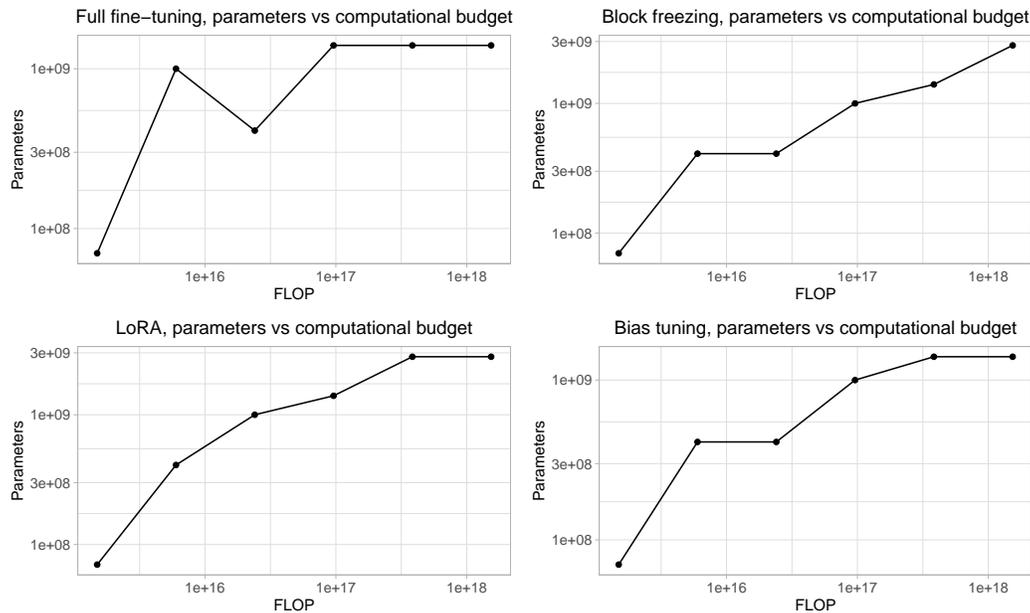


Figure 15: Optimal choices of model size given the computational budget, with respect to the performance on the downstream task evaluation.

For each of the task above, we specify its category in brackets. Since Tatoeba is not featured in the MTEB leaderboard,³ we skip it when calculating our approximation of the MTEB performance of the embedding models.

In principle, one could evaluate all the checkpoints on all the MTEB tasks. However, we decided to evaluate on a representative subset in order to speed up evaluation (evaluating some MTEB task takes more than 24 hours, and we have close to 1000 checkpoints, so the overall cost is significant).

C Correlation between loss and MTEB performance

In our work, we use loss as the main metric used to formulate the conclusions in form of the compute-optimal recipe and the scaling laws. This is in line with other works investigating scaling laws in various contexts, like [Hoffmann et al., 2022, Muennighoff et al., 2023a, Zhang et al., 2024]. However, we additionally evaluated the trained models on a subset of MTEB benchmark. We calculated the Spearman’s rank correlation between the loss and the MTEB performance and it is equal to -0.892 , which is significant. Additionally, Figure 16 visually compares the loss and MTEB performance.

D Data-Constrained Setup

In our work, we mainly focus on finding the compute-optimal recipe. The training datapoints we share can however be utilised to obtain conclusions about other setups. In particular, we can consider the data-constrained setup. Here, we compare models trained with the same amount of training data instead of the same amount of computation. In Figure 17 we present the plot the final contrastive loss vs. the amount of tokens seen by the network on a log-log scale. The conclusion that can be drawn is that in a data-constrained setup, one should use the largest model available and fine-tune it using full fine-tuning or LoRA with a rank on the bigger side.

E Hyperparameters for the training setup

Below, we list the hyperparameters that we used for training.

³<https://huggingface.co/spaces/mteb/leaderboard>

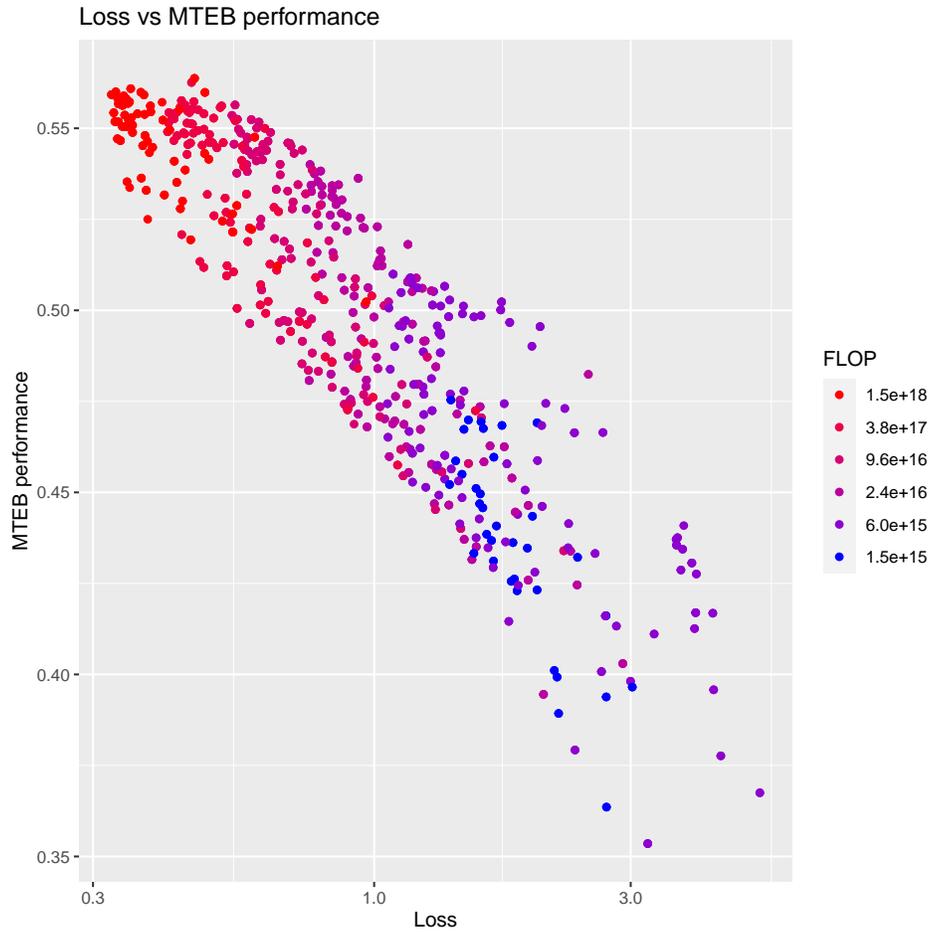


Figure 16: All the final checkpoints involved in our suite of experiments, evaluated both in terms of the final training loss and the performance on a representative subset of MTEB benchmark. The correlation between these two metrics is significant.

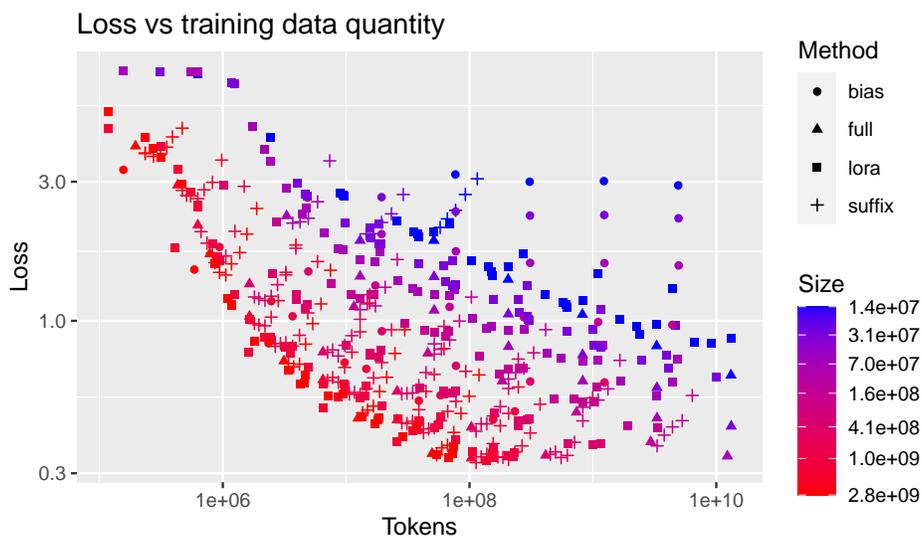


Figure 17: Final training loss vs number of training tokens for all the models. It is visible that given the number of training tokens, the best performance is obtained by the largest model.

```

batch_size = 1024
context_length = 75
AdamW.weight_decay = 0.1
tau = 0.025 # temperature parameter

```

E.1 Learning rate for block freezing and full fine-tuning

For all the experiments with block freezing and full fine-tuning, we base our peak learning rate on the peak learning rate from Pythia [Biderman et al., 2023]. We calculate our learning rate by dividing the Pythia learning rate by ten. We use the cosine learning rate schedule with warmup, warmup period being fixed to 10% of the training steps. Learning rate starts at zero, goes up to the peak learning rate, then follows the cosine schedule back to zero.

E.2 Temperature and weight decay for partial tuning

We found the temperature to affect the training performance rather significantly, while the weight decay was of a very slight importance. We run a grid search for partial tuning, varying temperature and weight decay and arrived at a choice of $\tau = \frac{1}{40}$ and weight decay 0.1. We compared the runs by calculating the results on a subset of MTEB benchmark. Grid results for temperature are detailed in Table 1, with best scores for each of the temperature value listed.

Table 1: A comparison of different temperature values

Temperature ⁻¹	Avg. score
50	0.5
40	0.51
35	0.5
30	0.5
10	0.47
1	0.3

E.3 Readout: last vs average

We found that using the averaging of embeddings as the pooling method results in a better performance over selecting the last embedding. We compare the results of the best run for mean and last pooling on full MTEB without MSMarco subset. The scores are presented in Table 2.

Table 2: Mean pooling works better than last pooling.

Pooling method	Avg. score
Mean pooling	0.47
Last pooling	0.36

E.4 LoRA hyperparameters

In order to be able to compare the loss from LoRA with partial block freezing, we leave most of the hyperparameters unchanged. We change the learning rate, we find the optimal learning rate by running a grid search over hyperparameters. The result shows that for 3 different rank choices, the same learning rate works best. The exact values are presented in Table 3.

E.5 Bias tuning hyperparameters

Similarly, as with LoRA, for bias tuning to be comparable with partial block freezing, we leave most of the hyperparameters the same. We vary the learning rate, we find the optimal learning rate by running a grid search over hyperparameters. The exact values can be seen in Table 4.

Table 3: Learning rate for LoRA.

Learning rate	LoRA rank	Loss
1e2	8	1.96
1e3		1.35
1e4		2.22
1e2	16	1.96
1e3		1.34
1e4		2.22
1e2	32	2.09
1e3		1.32
1e4		2.19

Table 4: Learning rate for bias tuning.

Learning rate	Loss
1e2	1.24
1e3	1.37
1e4	2.97

F Scaling laws

We aim to express the formula describing the behaviour of the loss for our models as a function of the number of parameters N and training tokens D . Following Hoffmann et al. [2022] we start with:

$$L(N, D) = C + \frac{A}{N^\alpha} + \frac{B}{D^\beta},$$

where $A, B, C, \alpha, \beta \in \mathbb{R}$ are parameters to be fitted based on the data.

Our setting differs from that of Hoffmann et al. [2022] as we train using the contrastive objective instead of the next-token prediction task, and moreover we apply various compute-efficient methods making some fraction of the parameters of the model *not trainable*. Despite that, we can still observe certain linear trends that the loss follows. For all the methods, when the FLOP budget is not overly restricted, if we fix the amount of both *all* the parameters as well as the *trainable* parameters and vary the amount of FLOP, the loss behaves approximately linearly. We observe that for full fine-tuning, the points are roughly collinear, as in Figure 18. On the other hand, for LoRA and partial block freezing, we discover that while the lines for one model are parallel, they are translated, depending on the amount of trainable parameters, which can be observed in Figures 21 and 22.

Following the Frantar et al. [2023]’s approach, we replace the term $\frac{B}{D^\beta}$ by $\frac{(a_s(1-S))^{b_s} + c_s}{D^\beta}$, where S is an additional parameter indicating the fraction of the model’s parameters that are trainable, and a_s, b_s are real-valued constants. Similarly, if we keep the forward token amount fixed and vary the amount of model parameters, the linear trend can be observed. Again, the lines are parallel and differ by an intercept, that seems to be linearly dependent on the logarithm of forward tokens, shown in Figure 20. Therefore, we replace the term $\frac{A}{N^\alpha}$ by $\frac{a_d \cdot \log(D) + b_d}{N^\alpha}$, where a_d, b_d are real-valued constants. The final formula describing the loss in our setting has the form:

$$L(S, N, D) = C + \frac{a_d \cdot \log(D) + b_d}{N^\alpha} + \frac{a_s \cdot (1 - S)^{b_s} + c_s}{D^\beta}.$$

We split our dataset into train and test set, with test set being the points corresponding to Pythia 2.8B which is the biggest model that we consider and train set the rest. Similarly to Hoffmann et al. [2022], we use L-BGFS optimization algorithm and Huber loss. We find that $\delta = 0.001$ prevents over-fitting and decreasing it further does not bring more benefits. We then compare the base formula being fitted from Chinchilla to our modified formula and find that our works better, as demonstrated in Figure 19. Both formulas are fitted after a grid search on the initial values of their coefficients. It can

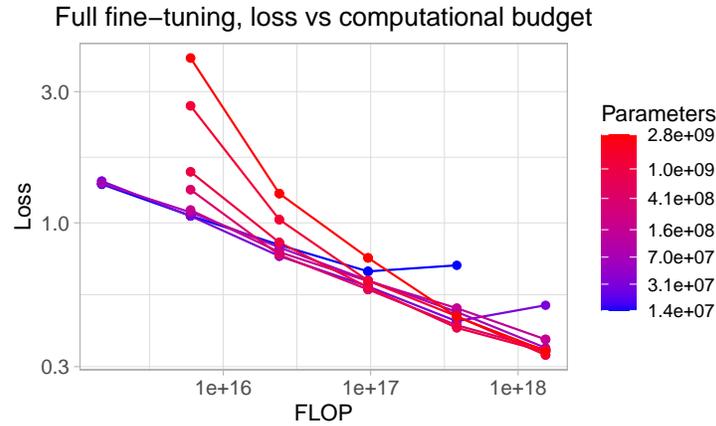


Figure 18: Loss on the y axis and FLOP on the x axis. Lines of the same colours correspond to the same amount of parameters in the model. Excluding points corresponding to severely constrained FLOP budget, the loss behaviour can be approximated by straight lines.

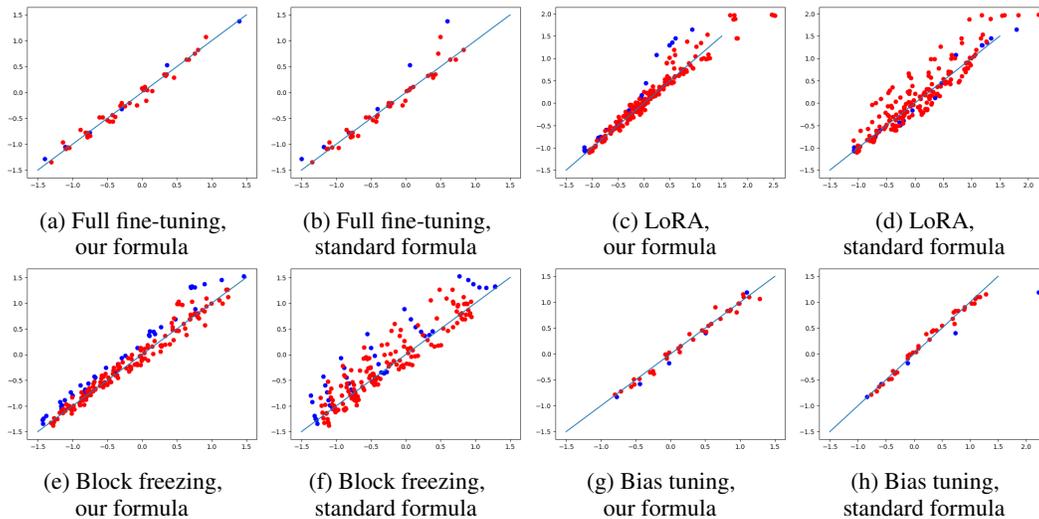


Figure 19: The actual loss on the y axis and predicted loss on the x axis. Red points are the training set and blue the test set. Figures 19a, 19c, 19e, and 19g are the results of fitting our formula for different fine-tuning methods, while Figures 19b, 19d, 19f, and 19h are the result of fitting the standard formula from Hoffmann et al. [2022]. It can be observed, that in contrast to the standard formula, our results in plots that are roughly linear, without additional trends visible.

be observed that we are able to accurately predict the loss for the bigger model for all the methods. Therefore, we hypothesize that our formula can be used for prediction the performance of models of an even larger scale.

G Robustness to increase in batch size and context length

Constrained by the computational budget, we performed the biggest line of experiments using batch size 1024 and context length 75. However, an increased context length as well as batch size might be desirable in many cases. Here, we demonstrate that the scaling laws established by us extend to a setting with batch size 2048 and context length 512. For those experiments, we trained five pre-trained decoder-only models from the Pythia suite [Biderman et al., 2023], with sizes 14M, 31M, 70M, 160M, and 410M. We consider five computational budgets: $1.5e15$, $6e15$, $2.4e16$, and $9.6e16$ FLOP. We fit the formula following the method from Appendix F, with the change of using the

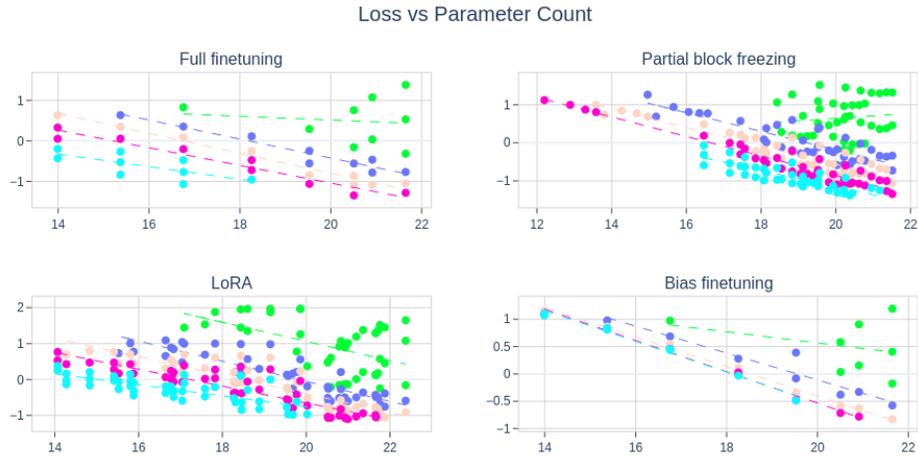


Figure 20: Logarithm of loss on the y axis and logarithm of the parameter count on the x axis. The colors represent different quantiles of the amount of tokens used in the model fine-tuning. Linear dependency can be observed, with the trends being approximately parallel for different quantiles.

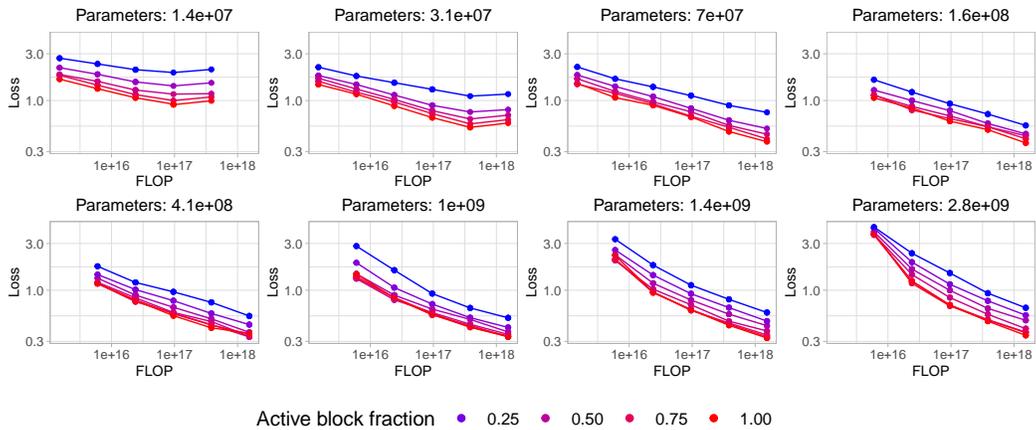


Figure 21: Plots of losses for block freezing. Each subplot contains datapoints corresponding to one model size. Loss on the y axis and FLOP on the x axis. Lines of the same colours correspond to the same amount of trainable parameters in the model. The loss behaviour can be approximated by straight lines.

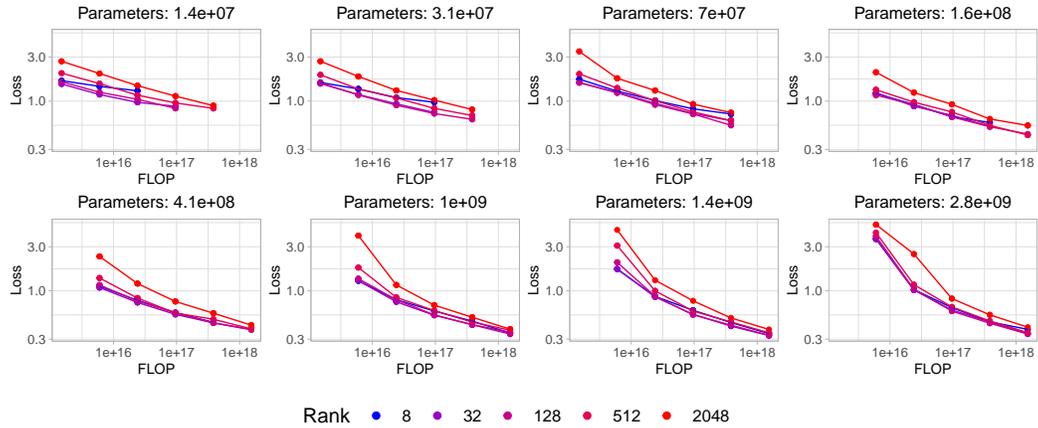


Figure 22: Plots of losses for LoRA. Each subplot contains datapoints corresponding to one model size. Loss on the y axis and FLOP on the x axis. Lines of the same colours correspond to the same amount of trainable parameters in the model. The loss behaviour can be approximated by straight lines.

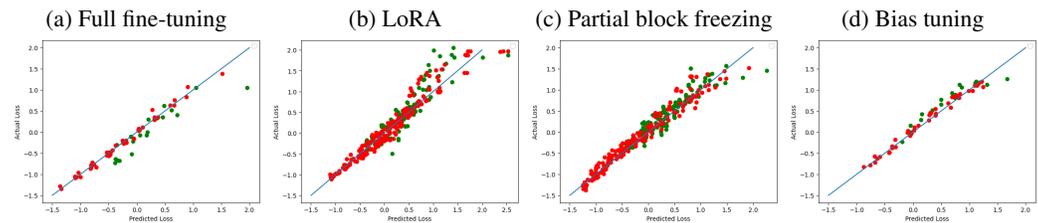


Figure 23: Plots of predicted vs actual losses for the 4 finetuning methods. The formula is fitted on the data coming from trainings with batch size 1024 and context length 75 (red points). It is then used to predict the losses for data points trained with batch size 2048 and context length 512 (green points).

number of training steps as the measure of data seen by the model instead of forward tokens. We then use it to predict the loss for the changed hyperparameters values. The results are visible in Figure 23. We conclude that our findings are relevant for setups with a larger batch size and context length.

H Computational resources

For all the trainings, we used a cluster with A100 GPUs with 40GB VRAM or with 80GB VRAM. We used 4 CPUs and 128GB of RAM per GPU. We estimate that the results contained in this article were obtained by expanding somewhere in between 30k and 40k GPU hours.

I Dataset

We utilize the BAAI dataset available at <https://data.baai.ac.cn/details/BAAI-MTP>.

It has a custom licence that allows for conducting academic research.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The main claims made in the abstract and introduction are an accurate reflection of the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The paper discusses the limitations extensively.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA] .

Justification: Our paper does not contain theoretical results, only empirical ones.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We fully disclose the information needed to reproduce the results in the paper. Since the reproduction of the training results would be rather costly, we also share the data we have got as the result of the training sweep, as an intermediate step in reproduction of conclusions/results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We will open-source the code used for training. We share the code as a supplementary material to the submission. We share the link for downloading the dataset in the zip file with the code. We moreover share the training results used for creating the plots here.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We specify all the training and test details. We moreover share the code that can be used to check those details without the necessity for consulting the paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Running the experiments several times would be too costly. However, we would like to point out that it is known that language modelling generally results in stable trainings.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide the information on the computational resources needed for the experiments in Appendix H.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The research does conform to the NeurIPS Code of Ethics in every respect.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: As this work aims at improving the efficiency for people to obtain embedding models, we only see positive social impacts. This was discussed in Section 6.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite the original paper that has produced the dataset we use. We provide the dataset webpage, which contains the licence in Appendix I. We cite the relevant packages that we used in our code in section 4.1.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: The main asset introduced in the paper is the code, that is attached as a supplementary material and it does contain the documentation.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.