
Don't Sweep your Learning Rate under the Rug: A Closer Look at Cross-modal Transfer of Pretrained Transformers

Danielle Rothermel¹ Margaret Li² Tim Rocktäschel^{1,3} Jakob Foerster¹

Abstract

Self-supervised pre-training of large-scale transformer models on text corpora followed by fine-tuning has achieved state-of-the-art on a number of natural language processing tasks. Recently, Lu et al. (2021) claimed that *frozen* pretrained transformers (FPTs) match or outperform training from scratch as well as unfrozen (fine-tuned) pretrained transformers in a set of transfer tasks to *other modalities*. In our work, we find that this result is, in fact, an artefact of not tuning the learning rates. After carefully redesigning the empirical setup, we find that when tuning learning rates properly, pretrained transformers do outperform or match training from scratch in all of our tasks, but only as long as the *entire model* is fine-tuned. Thus, while transfer from pre-trained language models to other modalities does indeed provide gains and hints at exciting possibilities for future work, properly tuning hyperparameters is important for arriving at robust findings.

1. Introduction

Transformer-based pretrained language models (LMs) have led to a revolution in the area of natural language processing (NLP) in recent years (Vaswani et al., 2017), a progress fueled by larger training data sets, bigger models and increasing computational demand. Finetuning such pretrained LMs has led to state-of-the-art performance across a variety of NLP benchmarks (Petroni et al., 2021; Wang et al., 2018; 2019). In these settings, the LM is pretrained on a large collection of natural language texts such as the Google Book Corpus (Zhu et al., 2015) or the 1B Word Benchmark (Chelba et al., 2014). Subsequently, the model is fine-tuned on a given task of interest, e.g. sentiment analysis (Maas




Figure 1. Test accuracy on the CIFAR10 LRA task across the learning rate sweep, with error bounds across 3 seeds. The learning rate reported by Lu et al. (2021), 1×10^{-3} , is marked with a dashed red line, demonstrating that any lower learning rate would have given inverted results on this task.

et al., 2011) or text classification (Zhang et al., 2015).

While the ability to transfer language representations, e.g. word representations (Mikolov et al., 2013; Pennington et al., 2014) or contextual representations (Devlin et al., 2019; Radford et al., 2019), between different language tasks has been well studied and revealed few-shot (Brown et al., 2020) and zero-shot (Petroni et al., 2019) abilities, recent work has focused on the exciting possibility of transfer between different modalities (Lu et al., 2021). Successful transfer between different modalities (e.g. natural language to images) can be interpreted as less of a pretraining of transferable representations but instead transfer of the general *computational structure* inherent in language to other tasks (Lu et al., 2021). Indeed, Lu et al. find that finetuning only the input and output layers of a fully trained NLP transformer model, *frozen* pretrained transformers (FPTs), matches or outperforms training from scratch of the same model, across a variety of tasks in different modalities.

In this paper, we report that the performance gains of FPTs disappear under fair tuning of the learning rate and one of the main claims of the paper no longer holds. FPTs perform better than random frozen models but are significantly worse than training from scratch on all four tasks studied. Concretely, with a lower learning rate, the ordering of the performance between the unfrozen and frozen variants is

¹Facebook AI Research ²University of Washington ³University College London. Correspondence to: Danielle Rothermel <drotherm@fb.com>.




Figure 2. [Left] Test accuracy of the best learning rate for all settings and tasks, with error bounds over 3 seeds. Exact values reported in Table 1. The Frozen variants under-perform across all tasks, but the Unfrozen Pretrained variant matches or exceeds all other variants. [Right] The test accuracy on the ListOps task across the learning rate sweep, error bounds over 3 seeds. The learning rate reported by Lu et al. (2021), 1×10^{-3} , is marked with a dashed red line, demonstrating that any lower learning rate would have given inverted results on this task. Each learning rate evaluated between 1×10^{-5} and 1×10^{-3} leads to a different conclusion about the best architecture.

inverted for all tasks (see Figure 1). The impact of hyperparameters on the empirical results in ML papers has been the subject of intense debates. It is not uncommon for careful finetuning to change the outcome of a paper or even undo years of “progress” in the field. For example, Melis et al. (2018) found that when fairly optimizing for hyperparameters, vanilla LSTM match more sophisticated recently introduced RNN variants across a variety of NLP tasks.

That said, interestingly, we find that when *not* frozen, Transformers do provide gains through transfer from text to other modalities. In particular in the challenging CIFAR10-LRA task (Tay et al., 2021), which consists of sequentialized CIFAR images, finetuning the entire pretrained model outperforms training from scratch by a large margin. On MNIST (LeCun & Cortes, 2010), the gap is small but significant and on the other two tasks finetuning the pretrained model matches the performance of training from scratch. This opens up exciting avenues for future research and we hope that our work will help the community avoid some potential pitfalls around hyperparameter tuning in the pursuit of this work, ensuring that the findings will stand the test of time.

2. Problem Setting and Background

The recent work from Lu et al. (2021) investigates the capability of transformers, pretrained on natural language, to generalize to other modalities with minimal finetuning. They limit the finetuning by freezing the majority of the weights in the residual layers, and report that this Frozen Pretrained Transformer (FPT) architecture achieves comparable or better results than transformers trained from scratch across their chosen tasks.

Lu et al. consider classification tasks across a range of modalities, including bit manipulation (Miconi et al., 2018), equation evaluation (Tay et al., 2021; Nangia & Bowman,

2018), image classification (Krizhevsky, 2012) and protein classification (Rao et al., 2019; Fox et al., 2013; Hou et al., 2018). The input is chunked into tokens for processing with the GPT2 architecture, with the tokens for the vision tasks being a flattened representation of a 4 by 4 pixel patch. The authors also include a more challenging version of the CIFAR10 task, CIFAR-LRA (Tay et al., 2021) where the patches are a single pixel, resulting in longer sequences.

FPTs as proposed by Lu et al. have the feedforward and multi-head attention frozen in each of the residual blocks. Only the input and output layers, layer norm parameters, and positional embeddings are finetuned. The authors compare this performance with a Frozen Random transformer and an Unfrozen variant. For the Unfrozen variant they report numbers from different architectures for each task we consider. For CIFAR10-LRA and ListOps they report numbers from a vanilla Transformer with tuned hyperparameters as provided in Tay et al. (2021). For MNIST and CIFAR10 they report results from GPT2, with CIFAR10 using a 3 layer model due to instability in training the full sized version.

The authors report training on a single learning rate (1×10^{-3}) for all tasks except Homology, and appear to report the results from a single seed per variation. The released code¹ along with the paper does not use a validation set and they report the test accuracies from a held-out test set.

3. Methods

Training of deep neural networks can be highly sensitive to the learning rates used for optimizing the network (Choi et al., 2020). Therefore, a natural question is to ask whether the results reported in Lu et al. (2021) have been impacted by the choice of using a fixed learning rate. To investigate

¹<https://github.com/kzl/universal-computation>

this, we rerun the experiments of Lu et al. while broadly sweeping the learning rate. As we will see later, any given *fixed* learning rate greatly changes the results.

To investigate the effects of pretraining and freezing across tasks from different modalities, we evaluate on four of the tasks explored by Lu et al.: ListOps, MNIST, CIFAR10 and CIFAR10-LRA. We do not replicate the Bit tasks because the transformers were able to perfectly solve them in the original work. The Homology task is not supported in the released codebase so it would be more difficult to ensure an accurate reproduction of their experimental setting. We evaluate the performance on the base GPT-2 model, at 12 layers. As in their work, we experiment with using transformer models pretrained on natural language, and with freezing the self-attention and feedforward layers finetuning only the input and output layers, the layer norm and the positional embeddings. Specifically, we consider:

Frozen Pretrained: The Frozen Pretrained Transformer introduced by Lu et al.

Frozen Random: The transformer is randomly initialized and the self-attention and feedforward layers are frozen before finetuning.

Unfrozen Pretrained: The transformer is initialized with a pretrained language model and finetuned without freezing any layers.

Unfrozen Random: The transformer is randomly initialized and finetuned without freezing any layers.

For each of these settings and tasks, we train using the Adam optimizer (Kingma & Ba, 2015) and sweep the learning rate logarithmically from 1×10^{-6} to 0.01. We use a batch size of eight and train up to a predefined maximum number of gradient steps (see Appendix 5 for details). We determine the training step for early stopping based on the performance on the validation set to obtain the best model from each run and to identify the best learning rate across the sweep. For each setting, we repeat the experiments with three seeds and report the mean test accuracy along with the standard error of the mean as measured on the held-out test set.

For the ListOps task the validation split is provided as part of the dataset released with the Long Range Arena (Tay et al., 2021), but for the CIFAR10 and MNIST datasets we create our own validation set. The 50K image train split in the CIFAR10 dataset is further subdivided into 5 training batches of 10K images, each containing a balanced number of samples from each class. We select one of these batches to be the validation set and train on the other 40K images. For MNIST, we split the 60K image training set into a 50K image training split and a 10K image validation set by randomly selecting 1000 images from each class.

For the ListOps task we note that the codebase released by

Lu et al. is affected by issues only recently fixed in the Long Range Arena codebase.² Specifically, the ListOps dataset includes sequences ranging from 500 to 2000 tokens, but the dataset tokenization utility truncated all sequences to 512 tokens. In addition, the “close” parenthesis for the list of operations was not tokenized due to not being an alphanumeric character. Between these two issues it was impossible to solve the long sequences of operations provided by the dataset. We resolved these issues in the dataset tokenization utility and adapted the architecture choices accordingly, by increasing the context length and the number of positions used by the transformer architecture to 2000. This change is possible because in all settings we fine-tune the positional embeddings. The additional tokenized character increased the token dimension for the ListOps task to 16.

Lu et al. report that the model capacity impacts the performance of each of the settings, with increased model capacity hurting the performance of the Unfrozen variants and helping the performance of the Frozen variants, resulting in some of the Unfrozen results being reported for models with 3 layers instead of the full 12 layer GPT2. To evaluate the impact of model capacity and to provide a datapoint between the two model sizes, we also test using a pretrained DistilGPT2 which is a 6 layer transformer distilled from the full sized pretrained GPT2 model.

4. Results and Discussion

While at a high level, our work confirms the finding from Lu et al. that transfer from NLP tasks to other modalities is indeed possible through finetuning, our results contradict theirs regarding which parts of the model should be fine-tuned. Our main finding is that while the *Unfrozen* Pretrained variant matches or outperforms all other settings across all tasks explored, the *Frozen* variants often greatly lag in performance comparatively, in direct contradiction to their findings. Table 1 compares between the different settings for each of the tasks across 3 seeds. For each task, the Frozen Pretrained setting outperforms the Frozen Random setting. However, in contrast to their results, the Unfrozen variants always outperform the Frozen variants and by a large margin for all the tasks except for MNIST.

The differences between these results and the ones obtained and reported by Lu et al. in their Section 3.1, 3.2 and 3.11 can be explained by investigating the test accuracy across the learning rate sweep, shown in Figure 1 for the CIFAR10-LRA task. Notably, the learning rate impacts not only the absolute performance but also the ordering between the settings. The learning rate reported by Lu et al., 1×10^{-3} , is marked with a vertical dashed red line. Since 1×10^{-3} is just before a precipitous drop in the performance of the unfrozen

²<https://github.com/google-research/long-range-arena>

Table 1. Comparison of test accuracy across initialization and finetuning methods for the GPT2 architecture.

	LISTOps	MNIST	CIFAR10	CIFAR10-LRA
FROZEN RANDOM	38.9 ± 0.3	98.0 ± 0.0	61.8 ± 0.2	44.2 ± 0.3
FROZEN PRETRAINED	46.1 ± 0.3	98.5 ± 0.1	66.3 ± 0.0	54.7 ± 1.4
UNFROZEN RANDOM	57.6 ± 0.8	98.7 ± 0.0	77.8 ± 0.2	62.0 ± 0.7
UNFROZEN PRETRAINED	56.3 ± 0.9	99.0 ± 0.0	77.7 ± 0.1	67.8 ± 0.3

transformers, had the authors picked a lower learning rate they would have arrived at very different conclusions.

When repeating this analysis for the ListOps task, in Figure 2, we see an even greater dependence on the learning rate. Each of the LRs evaluated between 1×10^{-5} and 1×10^{-3} results in different orderings and, hence, conclusions about the optimal architecture variant. See Appendix A for plots for MNIST and CIFAR10 tasks which uphold these findings.

The key shared finding between Lu et al. and our work is the benefit of finetuning from a model pretrained on natural language, even for tasks of different modalities. In their Section 3.2, Lu et al. find that the Frozen Pretrained transformer is superior to a Frozen Random variant. We verify that pretraining improves performance across all tasks for the Frozen transformers, and in addition find that for some tasks pretraining provides benefits for finetuning the Unfrozen variants. For the CIFAR10-LRA task, the Unfrozen Pretrained variant outperforms all other variants by 4.8%, and on MNIST the Unfrozen Pretrained variant outperforms the rest by a small margin. This benefit from pretraining on some tasks, paired with matched performance on the rest, suggests that it may be expedient to run initial experiments in new settings by finetuning from a natural language pretrained model. However the varying success of pretraining across tasks raises an open question, for future work, about which qualities of a task lead to benefits from pretraining.

In their Section 3.4, Lu et al. compare the computation efficiency between the Frozen Random and Frozen Pretrained variants by reporting the number of gradient steps to convergence. When sweeping the learning rate, we see that the final performance of the Frozen variants is much lower than the Unfrozen variants. Thus, we instead compare computational efficiency by reporting the number of gradient steps to match the performance of the Frozen Pretrained variant in Appendix B. The Frozen Random variant does not match performance in any of the tasks, verifying the authors’ assertion that pretraining improves the computational efficiency of the Frozen variants. However, for all tasks the Unfrozen variants require fewer gradient steps. For all but ListOps the Unfrozen Pretrained variant requires the least gradient steps, demonstrating that in some tasks pretraining helps not only final performance but also computational efficiency.

In their Section 3.6, Lu et al. report that the Frozen Pre-

trained models underfit the CIFAR10 task. We validate these findings (in Appendix C) but argue that underfitting is likely what leads to the poor comparative performance of the Frozen variants in our experiments. In addition, while the Frozen variants underfit on the MNIST and CIFAR10 tasks, the Frozen Pretrained variant has the largest train/test gap of all settings on ListOps, invalidating the hypothesis that the Frozen variants always underfit the data.

In their Section 3.7, Lu et al. report that increasing the model capacity of the Frozen Pretrained setting improved the performance on the CIFAR10 task, suggesting an easy way to achieve performance gains. We report results from similar experiments, with the addition of the learning rate sweep, in Appendix D, confirming their finding of performance gains from increased capacity for the Frozen Pretrained variant on CIFAR10. Our results also verify that these claims hold for the CIFAR10-LRA task and but show that increasing model capacity also benefits the Unfrozen Pretrained variant. Interestingly, the increase in model capacity improves the Unfrozen Pretrained setting across all tasks whereas it only improves the Unfrozen Random setting on CIFAR10-LRA.

In their Section 3.11, Lu et al. describe the impact of unfreezing part or all of the pretrained model. They report that unfreezing both the feedforward layers and the attention heads is detrimental to performance. This experiment corresponds to our Unfrozen Pretrained variant which we find outperforms all other variants when the learning rate is properly swept, contrary to their findings.

5. Conclusion

Transformer architectures pretrained on natural language texts are some of the most successful models in NLP in recent years. Therefore, investigating how they can best be used as a starting point for solving tasks in other modalities is an important research direction. In our work, we show that, across a variety of tasks, the best results are obtained when finetuning all of the weights of a pretrained model. These results directly contradict prior work which concluded that freezing most of the model leads to superior performance. We demonstrate that these prior conclusions were an artefact of using a specific, fixed learning rate, and hope that our work will help pave the way for robust investigations into cross-modal transfer in the future.

6. Acknowledgements

The authors wish to thank Edward Grefenstette, Hengyuan Hu and Patrick Lewis for many helpful discussions during the initial stages of this project.

References

- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pp. 1877–1901. Curran Associates, Inc., 2020.
- Chelba, C., Mikolov, T., Schuster, M., Ge, Q., Brants, T., Koehn, P., and Robinson, T. One billion word benchmark for measuring progress in statistical language modeling. In *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- Choi, D., Shallue, C. J., Nado, Z., Lee, J., Maddison, C. J., and Dahl, G. E. On empirical comparisons of optimizers for deep learning. In *International Conference on Learning Representations*, 2020.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423.
- Fox, N. K., Brenner, S. E., and Chandonia, J.-M. Scope: Structural classification of proteins—extended, integrating scop and astral data and classification of new structures. In *Nucleic Acids Research*, pp. D304–D309, 2013. doi: 10.1093/nar/gkt1240.
- Hou, J., Adhikari, B., and Cheng, J. Deepspf: Deep convolutional neural network for mapping protein sequences to folds. In *Bioinformatics*, pp. 1295–1303, 08 2018. doi: 10.1093/bioinformatics/btx780.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- Krizhevsky, A. Learning multiple layers of features from tiny images. *University of Toronto*, 05 2012.
- LeCun, Y. and Cortes, C. MNIST handwritten digit database. 2010. URL <http://yann.lecun.com/exdb/mnist/>.
- Lu, K., A. Grover, A., Abbeel, P., and Mordatch, I. Pre-trained transformers as universal computation engines. *arXiv:2103.05247*, 2021.
- Maas, A. L., Daly, R. E., Pham, P. T., Huang, D., Ng, A. Y., and Potts, C. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics.
- Melis, G., Dyer, C., and Blunsom, P. On the state of the art of evaluation in neural language models. In *International Conference on Learning Representations*, 2018.
- Miconi, T., Stanley, K., and Clune, J. Differentiable plasticity: training plastic neural networks with backpropagation. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3559–3568. PMLR, 10–15 Jul 2018.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. Distributed representations of words and phrases and their compositionality. In Burges, C. J. C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.
- Nangia, N. and Bowman, S. ListOps: A diagnostic dataset for latent tree learning. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pp. 92–99, New Orleans, Louisiana, USA, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-4013.
- Pennington, J., Socher, R., and Manning, C. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1532–1543, Doha, Qatar, October 2014. Association for Computational Linguistics. doi: 10.3115/v1/D14-1162.
- Petroni, F., Rocktäschel, T., Riedel, S., Lewis, P., Bakhtin, A., Wu, Y., and Miller, A. Language models as knowledge bases? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 2463–2473, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1250.

- Petroni, F., Piktus, A., Fan, A., Lewis, P., Yazdani, M., De Cao, N., Thorne, J., Jernite, Y., Karpukhin, V., Mail-lard, J., Plachouras, V., Rocktäschel, T., and Riedel, S. KILT: a benchmark for knowledge intensive language tasks. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2523–2544, 2021.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Rao, R., Bhattacharya, N., Thomas, N., Duan, Y., Chen, P., Canny, J., Abbeel, P., and Song, Y. Evaluating protein transfer learning with tape. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Tay, Y., Dehghani, M., Abnar, S., Shen, Y., Bahri, D., Pham, P., Rao, J., Yang, L., Ruder, S., and Metzler, D. Long range arena : A benchmark for efficient transformers. In *International Conference on Learning Representations*, 2021.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, u., and Polosukhin, I. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, pp. 6000–6010, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pp. 353–355, Brussels, Belgium, November 2018. Association for Computational Linguistics. doi: 10.18653/v1/W18-5446.
- Wang, A., Pruksachatkun, Y., Nangia, N., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. Super-glue: A stickier benchmark for general-purpose language understanding systems. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Le Scao, T., Gugger, S., Drame, M., Lhoest, Q., and Rush, A. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45, Online, October 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-demos.6.
- Zhang, X., Zhao, J., and LeCun, Y. Character-level convolutional networks for text classification. In Cortes, C., Lawrence, N., Lee, D., Sugiyama, M., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.
- Zhu, Y., Kiros, R., Zemel, R., Salakhutdinov, R., Urta-sun, R., Torralba, A., and Fidler, S. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV), ICCV '15*, pp. 19–27, USA, 2015. IEEE Computer Society. ISBN 9781467383912. doi: 10.1109/ICCV.2015.11.

A. Learning Rate Sweeps

We report the impact of the learning rate on the test accuracy for all tasks, MNIST, ListOps and CIFAR10 and CIFAR10-LRA. In all cases, the learning rate reported by Lu et al. (2021) is marked with a dashed red line and it is clear that reducing the reported learning rate would invert the findings.




Figure 3. Test accuracy of all tasks across the learning rate sweep, error bars across 3 seeds.

B. Computational Efficiency

We investigate the impact of pretraining and freezing on the computational efficiency of the architectures. Because the final performance varies dramatically between variants we compare the number of gradient steps necessary to reach the best test accuracy of the Frozen Pretrained variant. We find that the Frozen Random variant is never able to match this performance. The Unfrozen settings require fewer gradient steps to match performance across all tasks, with pretraining generally improving computational efficiency in three out of four tasks.

Table 2. The impact of pretraining on compute efficiency, comparing the number of gradient steps, per variant, to match the reported best mean test accuracy of the Frozen Pretrained variant.

	LISTOPS	MNIST	CIFAR10	CIFAR10-LRA
UNFROZEN RANDOM	1.0×10^5	1.1×10^5	7.0×10^4	1.8×10^5
FROZEN RANDOM	-	-	-	-
UNFROZEN PRETRAINED	1.9×10^5	4.6×10^4	4.3×10^4	5.3×10^4
FROZEN PRETRAINED	1.9×10^5	2.4×10^5	3.8×10^5	2.4×10^5

C. Underfitting versus Overfitting

We investigate the extent to which each of the variants is able to fit the data by reporting the train and test accuracy at the threshold where we ended training, specified in Table 5 for each task. We find that across the MNIST and CIFAR10 tasks the Frozen variants underfit the data. However, this trend does not hold for the ListOps task where the Frozen Pretrained setting has the largest train/test gap.

Table 3. Train versus Test accuracy at the maximum number of gradient steps taken for each task as listed in Table 5.

		LISTOPS	MNIST	CIFAR10	CIFAR10-LRA
UNFROZEN RANDOM	TRAIN	58.1	99.9	96.4	74.1
	TEST	57.5	98.6	77.7	61.7
	DIFF	0.6	1.3	18.7	12.4
FROZEN RANDOM	TRAIN	39.7	98.8	62.9	44.9
	TEST	39.2	98.0	61.1	43.8
	DIFF	0.5	0.8	1.9	1.1
UNFROZEN PRETRAINED	TRAIN	57.8	100.0	98.7	91.0
	TEST	55.7	99.0	77.1	67.0
	DIFF	2.1	1.0	21.7	24.0
FROZEN PRETRAINED	TRAIN	52.2	99.3	67.8	55.6
	TEST	46.4	98.5	65.5	53.4
	DIFF	5.8	0.8	2.3	2.2

D. Scaling Model Capacity

We investigate the impact of scaling the model capacity across three of the tasks, MNIST, CIFAR10 and CIFAR10-LRA. We compare the DistilGPT2 model at 6 layers against the GPT2 base model at 12 layers, both provided by the HuggingFace Transformers library (Wolf et al., 2020). Scaling has little or no impact on MNIST and the only variant to show improvement across all tasks with increased model capacity is the Unfrozen Pretrained setting. The Frozen Pretrained setting also improves with model capacity on both CIFAR10 tasks.

Table 4. The impact of scaling the size of the transformers across three of the tasks, comparing the performance of the DistilGPT2 architecture with that of the GPT2 architecture.

		MNIST	CIFAR10	CIFAR10-LRA
FROZEN RANDOM	DISTILGPT2	98.0 ± 0.1	60.1 ± 0.1	45.0 ± 0.1
	GPT2	98.0 ± 0.0	61.8 ± 0.2	44.2 ± 0.3
FROZEN PRETRAINED	DISTILGPT2	98.5 ± 0.1	65.2 ± 0.5	51.1 ± 0.4
	GPT2	98.5 ± 0.1	66.3 ± 0.0	54.7 ± 1.4
UNFROZEN RANDOM	DISTILGPT2	98.6 ± 0.1	77.5 ± 0.1	59.7 ± 0.2
	GPT2	98.7 ± 0.0	77.8 ± 0.2	62.0 ± 0.7
UNFROZEN PRETRAINED	DISTILGPT2	98.9 ± 0.0	76.8 ± 0.1	65.5 ± 0.5
	GPT2	99.0 ± 0.0	77.7 ± 0.1	67.8 ± 0.3

E. Evaluation Setup

We trained each of the tasks for a logarithmic sweep of learning rates, from 1×10^{-6} to 1×10^{-2} . Each task was run for a fixed number of gradient steps, specified in Table 5. The validation accuracy was used to perform early stopping and to identify the model in each run to evaluate and the test accuracy from that model is reported.

Table 5. Threshold number of gradient steps used to report test accuracy results, per task and model type.

TASK	MODEL TYPE	NUMBER GRADIENT STEPS
LISTOPS	GPT2	3×10^5
MNIST	GPT2	4×10^5
CIFAR10	GPT2	4×10^5
CIFAR10 LRA	GPT2	3×10^5