

Prefix-Propagation: Parameter-Efficient Tuning for Long Sequences

Jonathan Li^{2*}, Will Aitken^{1,2}, Rohan Bhambhoria^{1,2}, Xiaodan Zhu^{1,2†}

¹Department of Electrical and Computer Engineering, Queen’s University

²Ingenuity Labs Research Institute, Queen’s University

{jxl, will.aitken, r.bhambhoria, xiaodan.zhu}@queensu.ca

Abstract

Parameter-efficient tuning aims to mitigate the large memory requirements of adapting pretrained language models for downstream tasks. For example, one popular method, prefix-tuning (Li and Liang, 2021; Liu et al., 2022), prepends trainable tokens to sequences while freezing the rest of the model’s parameters. Although such models attain comparable performance with fine-tuning when applied to sequences with short to moderate lengths, we show their inferior performance when modelling long sequences. To bridge this gap, we propose *prefix-propagation*, a simple but effective approach that conditions prefixes on previous hidden states. We empirically demonstrate that prefix-propagation outperforms prefix-tuning across long-document tasks, while using $\sim 50\%$ fewer parameters. To further investigate the proposed architecture, we also show its advantage in calibration, and perform additional study on its relationship with kernel attention. To the best of our knowledge, this work is the first to focus on parameter-efficient learning for long-sequence language tasks.¹

1 Introduction

The Transformer architecture (Vaswani et al., 2017) has changed the landscape of recent natural language processing approaches by enabling the pre-training of state-of-the-art large language models (LLM) (Devlin et al., 2019; He et al., 2020; Brown et al., 2020). However, fine-tuning and storing full copies of LLMs can consume prohibitively large quantities of resources. Parameter-efficient fine-tuning (PEFT) methods such as prefix-tuning (Li and Liang, 2021; He et al., 2021a; Liu et al., 2022) address these concerns by reducing the number

Method	20-newsgroups	Hyperpartisan
Prefix-Tuning	69.7	75.3
Fine-Tuning	72.3	81.5

Table 1: Mean F1-Scores of prefix-tuning and fine-tuning Longformer for common long-document classification tasks.

of trainable parameters. Prefix-tuning can tune 0.01% of parameters and still match the performance of regular fine-tuning (updating all model parameters).

PEFT has been investigated for tasks with inputs consisting of sentences, sentence-pair, or sequences that fit within the typical LLM maximum tokens. However, the performance of PEFT for tasks with longer textual sequences has been overlooked. In this work, we investigate this oversight and provide evidence suggesting that the gap between PEFT and regular fine-tuning is substantial when modelling long sequences. As shown in Table 1, prefix-tuning underperforms fine-tuning on long sequence classification tasks, Hyperpartisan (Kiesel et al., 2019) and 20-newsgroups (Lang, 1995), when used with the popular long-document model Longformer (Beltagy et al., 2020).

In this paper, we propose a simple and effective method, *prefix-propagation*, which consistently improves the performance of PEFT for long sequence models. Unlike prefix-tuning, prefix-propagation propagates the hidden states corresponding to prefixes through the attention computation. This allows for the prefixes hidden states to dynamically change as the input propagates through each layer. To further understand prefix propagation, we investigate the reliability of the model’s predictions by performing analyses on calibration. Lastly, we conduct study on prefix-based methods in terms of kernel attention to strengthen their theoretical value.

In summary, our contributions are as follows:

*Work done during a student internship at Ingenuity Labs.

†Corresponding author.

¹Our code is publicly available at <https://github.com/MonliH/prefix-propagation>

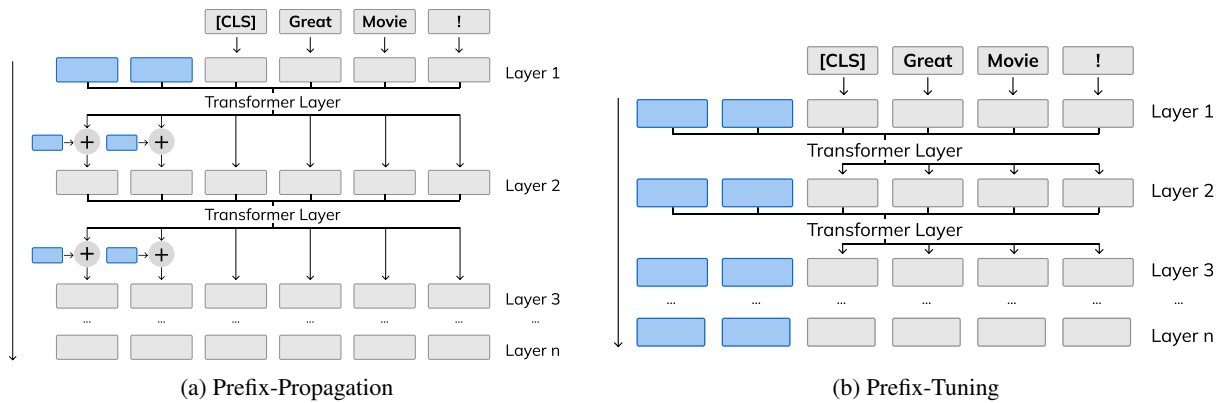


Figure 1: Illustration of the differences between (a) prefix-propagation (ours) (b) and prefix-tuning (Liu et al., 2022; Li and Liang, 2021). Blue blocks denote trainable prompts, and “Transformer Layer” represents the computation done in a layer of the pre-trained LLM. Note that in prefix-propagation (a), the summation of prefixes continues for layers beyond 3, up to n . This operation is encapsulated by the ellipses. In prefix-tuning (b), prefixes in subsequent layers do not depend on hidden states from past layers (they are simply overwritten).

- We study PEFT for long documents and show that prefix-tuning is significantly inferior to fine-tuning in this scenario. To the best of our knowledge, this is the first work to focus on PEFT for long documents.
- We introduce prefix-propagation, which consistently improves the performance over prefix tuning on the different long document datasets, while using 50% fewer parameters.
- We study the reliability of the predictions by performing analyses on calibration and show that models tuned with prefix-propagation are better calibrated.
- We elucidate the relationship between prefix-propagation and kernel attention and perform an ablation study that utilizes this insight.

2 Related Works

Long Sequence Models Numerous methods have been proposed to reduce the complexity of attention from $O(n^2)$ to $O(n)$ such as kernel approximations (Choromanski et al., 2020; Katharopoulos et al., 2020; Peng et al., 2021) and fixed (Child et al., 2019; Beltagy et al., 2020; Zaheer et al., 2020) or learned (Kitaev et al., 2020) sparse attention patterns. For a broader summary, please refer to Tay et al. (2022). In this work, we use Longformer (Beltagy et al., 2020). To linearize attention complexity, Longformer employs sliding window attention while globally attending to relatively few special tokens.

Parameter-Efficient Tuning Inspired by the success of manual prompting (Brown et al., 2020), prefix-tuning (Li and Liang, 2021; Liu et al., 2022) prepends trainable “soft” prompts to an input sequence. Although further PEFT methods have since been introduced (He et al., 2021a; Hu et al., 2021; Ben Zaken et al., 2022), we focus on adapting prefix-tuning. We note that our adaptation does not violate orthogonality and thus prefix-propagation can still be compounded with other PEFT methods as proposed in the UnifiedPET framework (He et al., 2021a), likely yielding similar performance gains. We leave the empirical validation of this hypothesis for future work.

Our work also adheres to the key motivation of the recent PEFT method, inducer-tuning (Chen et al., 2022), which is that optimal prefixes should be close to queries within their latent space. We derive queries, keys, and values from the same prefix token, limiting the distance that separates them.

3 Prefix Propagation

3.1 Methodology

In this section we introduce prefix-propagation, which, unlike prefix-tuning, propagates the hidden states corresponding to prefixes through the attention computation. This allows for the prefixes hidden states to dynamically change as the input propagates through each layer. Prefix-propagation and its predecessor, prefix-tuning are depicted in Figure 1a and 1b respectively. For the first layer of the transformer, we prepend j trainable prefixes (i.e.,

Method	% Tuned	WikiHop		ArXiv		20-newsgroups			Hyperpartisan		
		Acc	P	R	F1	P	R	F1	P	R	F1
ROBERTa PT	0.1	11.7	79.4	79.6	79.8	67.9	67.0	68.2	70.4	59.2	64.1
Prefix-Tuning	0.1	38.9	81.5	81.7	82.7	68.9	68.4	69.7	78.3	73.8	75.3
Prefix-Propagation	0.05	<u>42.2</u>	83.1	83.1	83.3	<u>70.1</u>	<u>69.7</u>	<u>71.0</u>	<u>86.4</u>	77.7	81.8
Fine-Tuning	100	74.0	83.1	<u>82.9</u>	83.3	71.8	71.2	72.3	87.8	<u>76.2</u>	<u>81.5</u>

Table 2: Main results of prefix-propagation compared to prefix-tuning and traditional fine-tuning on the validation sets of each dataset. All approaches use Longformer-base except “ROBERTa PT”, which is prefix-tuning on ROBERTa-base. Micro F1 and macro-average precision (“P”) and recall (“R”) is reported for ArXiv, Hyperpartisan (with mean across 5 runs), and 20-newsgroups. Accuracy is reported for WikiHop. Performance is reported on test splits with the exception of Hyperpartisan, which is performance on the validation split (See Appendix B for reasoning). The best run is **bold** and second best is underlined.

embeddings) to the input sequence (blue blocks in top left of Figure 1a). Then, before every subsequent layer, we sum new trainable matrices onto the first j embeddings corresponding to the prefixes (denoted by the sum operators in Figure 1a). By propagating instead of overwriting, we halve the number of parameters trained while simultaneously improving performance on long-document tasks.

We now formalize prefix-propagation. Multi-headed attention processes query, key, and value matrices derived from a sequence $C \in \mathbb{R}^{m \times d}$ with length m and embeddings of size d . Our method modifies traditional attention by concatenating a prefix $P \in \mathbb{R}^{j \times d}$ of length j to the sequence:

$$H_{l,i} = \text{Attn}(D^{(l)}W_q^{(l,i)}, \quad (1)$$

$$D^{(l)}W_k^{(l,i)}, D^{(l)}W_v^{(l,i)})$$

$$D^{(l)} = \begin{cases} \text{cat}(P^{(l)}, C) & \text{if } l = 1 \\ \text{cat}(P^{(l)} + C[:, j, :], C[:, j, :]) & \text{if } l > 1 \end{cases}$$

where inputs C are projected through pre-trained weight matrices $W_q^{(l,i)}, W_k^{(l,i)}, W_v^{(l,i)} \in \mathbb{R}^{d \times d_h}$ per layer l and head i yielding the output of the attention head, $H \in \mathbb{R}^{(j+m) \times d_h}$. The prefixes are concatenated for the first layer ($l = 1$) and summed to their corresponding hidden states for the remaining layers ($l > 1$). We do not continually concatenate new prefixes to the sequence to avoid increasing the sequence length after each layer.

For both prefix-tuning and prefix-propagation, prefixes (keys and values) are globally attended to by all queries. Unlike prefix-tuning however, our method concatenates additional hidden states *before* the hidden states C are projected by $W_k^{(i)}$ and $W_v^{(i)}$. By doing so, prefix-propagation modifies query matrices, allowing prefixes to attend to

other hidden states globally, thereby increasing representation capability. This approach is somewhat analogous to the external global tokens inserted in the BigBird-ETC model (Zaheer et al., 2020). By attending to other tokens, the prefixes can act as special storage tokens, which is particularly useful in the restricted regime of long-document modelling where relatively few tokens have global context. Conversely, prefix-tuning only concatenates trained key and value matrices, $P_k, P_v \in \mathbb{R}^{j \times d_h}$, statically to the sequence:

$$H_{l,i} = \text{Attn}(CW_q^{(l,i)}, \text{cat}(P_k^{(l,i)}, CW_k^{(l,i)}), \quad (2)$$

$$\text{cat}(P_v^{(l,i)}, CW_v^{(l,i)}))$$

Since our method has a single prefix matrix, P instead of separate P_k and P_v matrices, we reduce the number of trained parameters by 50%.

3.2 Calibration

We further study the proposed prefix-propagation method to understand the reliability of model’s predictions through calibration. Well-calibrated models output confidence scores that closely match the models’ accuracy. Either over-confident or under-confident models are undesirable. Calibration has widely been overlooked in PEFT methods. To quantify calibration in our work, we use expected calibration error (ECE), which bins predictions based on model confidence and compares them to accuracy (Pakdaman Naeini et al., 2015; Guo et al., 2017).

3.3 Kernel Decomposition

Traditional attention is analogous to applying a kernel smoother over inputs (Tsai et al., 2019).

Motivated by this insight, we reformulate prefix-propagation as a sum of kernelized attention modules. Separating the modules introduces flexibility in two ways: (1) Their individual kernel forms can be mixed and matched and (2) A hyperparameter scale factor α can be applied to the prefix component to increase or decrease its weighting. Equation 3 defines kernel decomposition for prefix-propagation²:

$$H = \text{Kern}(\text{cat}(P, C)W_q, CW_k, CW_v) + (\alpha)\text{Kern}(\text{cat}(P, C)W_q, PW_k, PW_v) \quad (3)$$

where Kern refers to kernel attention as formulated in (Tsai et al., 2019). The first term results from attending to the original sequence, C , and the second comes from attending to the prefixes, P . We provide the derivation of Equation 3 and the full definition of kernel attention in Appendix A.

Our main motivation for presenting prefix decomposition is to establish foundational knowledge and guide future research. Ergo, we restrict experiments in this initial presentation to using just the default exponential kernel (Appendix A).

4 Experiments and Results

Datasets We evaluate our approach on three long-document classification tasks: ArXiv (He et al., 2019), an 11-class classification task composed of academic research papers, the 20-newsgroups (Lang, 1995) classification task consisting of mailing lists that fall into one of 20 classes, and the Hyperpartisan dataset, a binary classification task for extremist news classification (Kiesel et al., 2019). We also run experiments on WikiHop (Welbl et al., 2018), a long-document reading comprehension task requiring multi-step reasoning.

Due to compute limitations inherent to working with long documents, with the exception of Hyperpartisan, we only report a single run for each task. This mimics the original Longformer reporting scheme (Beltagy et al., 2020). For Hyperpartisan, the smallest of the datasets, we report mean metrics averaged over five seeds.

Baselines As a baseline, we fine-tune Longformer-base (approx. 149M parameters) as closely as possible to Beltagy et al. (2020). For PEFT, we evaluate prefix-tuning on Longformer-base and RoBERTa-base (approx. 125M parameters) (Liu et al., 2019).

²We omit layer, l and head, i for brevity.

Method	ArXiv	HY.	NG.
ROBERTa PT	<u>0.056</u>	0.228	0.123
Prefix-Tuning	0.075	0.153	0.117
Prefix-Propagation	0.042	0.093	<u>0.122</u>
Fine-Tuning	0.099	<u>0.138</u>	0.212

Table 3: ECE scores of tested approaches. Lower is better. **Bold** is the best and underline is the second best. “HY.” is Hyperpartisan, and “NG.” is 20-newsgroups.

More details on dataset sizes, pre-processing, and hyperparameters are in Appendix B.

4.1 Results and Discussion

Across all tasks, our results in Table 2 verify that prefix-tuning is inferior to fine-tuning long sequences. Conversely, prefix-propagation consistently outperforms prefix-tuning and is comparable to fine-tuning on most tasks. Prefix propagation also performs competitively on Hyperpartisan, a relatively small dataset with only 625 samples. This is in contrast to prefix-tuning, which is known to underperform in low-data settings (Gu et al., 2022). Because we ran multiple seeds on Hyperpartisan, we also found that prefix-propagation’s better performance relative to prefix-tuning is statistically significant ($p < 0.05$, using a single-tailed t-test). We do not have multiple samples to run these tests for larger datasets, but we emphasize that Hyperpartisan likely has the most variance and yet it is still statistically significant. We suspect that prefix-propagation’s performance exceeds prefix-tuning because propagated prefixes can transmit global context across multiple layers, possibly modelling more expressive abstractions.

We note one exception where prefix-based methods still leave room for improvement: multiple-choice question answering on WikiHop. We hypothesize that prefix methods have insufficient capacity to properly model complex long-document multi-step question answering.

We also observe that prefix-based methods, and especially prefix-propagation, achieve better calibration than fine-tuning, as shown in Table 3. Unlike prefix-tuning however, prefix-propagation effectively balances calibration with accuracy metrics. The calibration of fine-tuning deteriorates as training progresses (Figure 4 in Appendix C) and we speculate that this may be due to catastrophic

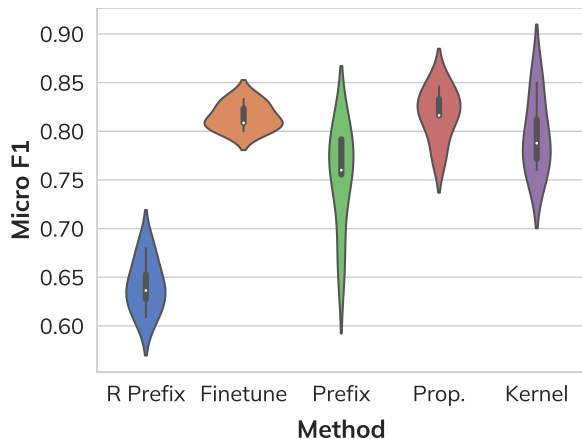


Figure 2: Violin plot of Micro F1 Score for five different seeds on the Hyperpartisan task. White dots, gray boxes, and gray lines are the medians, interquartile ranges, and ranges respectively. Width of the five violin shapes show the probability densities for the corresponding F1-score. All methods tune `Longformer`-base except “R Prefix”, which is prefix-tuning on `RoBERTa`-base.

forgetting (Jagielski et al., 2022).

As an initial test for our ongoing prefix-propagation kernel study, we show results on Hyperpartisan in Figure 2. The kernelized version of prefix-propagation achieves the best single-run performance, but has higher variance than fine-tuning and prefix-propagation which necessitates further research.

5 Conclusion

Our research focuses on parameter efficient tuning for long documents tasks. We introduce prefix-propagation, which consistently improves performance over prefix-tuning on long document datasets, while using 50% fewer parameters. We study the reliability of the predictions by performing analyses on calibration and show that models tuned with prefix-propagation are better calibrated. We lastly explicate prefix-propagation from a kernel perspective, uncovering insights for future PEFT research.

Limitations

Scope

This short paper serves as an initial step toward PEFT for long-document models. As such, our evaluated scope of models, tasks, datasets, and kernel variations is limited. We acknowledge the need to experiment across broader settings and hope our work provides a foundation for others to build on.

Future experiments should analyze the validity and efficacy of using prefix-propagation with other long-sequence models to determine whether the prefix modality is suitable for non-sparse attention approximations. For example, would the projection of prefix vectors using a random feature map as in Choromanski et al. (2020) result in an excessive loss of information for these critical tokens?

Regarding tasks and datasets, the performance degradation in prefix methods for WikiHop deserves significant attention. Verifying whether this extends to other reading comprehension and question-answering tasks will assist in guiding future research efforts. We restricted our research to the encoder-only version of `Longformer`, but using the encoder-decoder version, LED would enable analysis of sequence-to-sequence tasks. The SCROLLS benchmark (Shaham et al., 2022) would be a good starting point for this analysis since it includes an LED baseline.

Combining prefix and kernel methods is an ongoing research effort and there are several questions we plan to address: (1) What are the effects of swapping the default exponential kernel with other variants such as linear, polynomial, and RBF? (2) Does making the α scale parameter trainable improve performance? (3) Can we have a separate scale parameter for each query and should they be trainable? (4) Is this approach effective for modalities other than long-document? (5) Can we separate other components of attention into modular kernels (e.g. local and global kernels for sparse attention)?

Robustness

The size and nature of long-sequence tasks often resulted in long run times for the larger datasets ArXiv, 20-newsgroup and WikiHop. Consequently, we report results of one seed after doing a hyperparameter search for learning rate. This aligns with the reporting system of the original `Longformer` paper (Beltagy et al., 2020) but greater assurance in all long-sequence task performance could be achieved by accumulating results over several seeds. The size of datasets and iteration over several epochs somewhat mitigate this concern.

Ethics Statement

Our work helps to address the environmental and equitable distribution concerns of LLMs (Strubell et al., 2019). All PEFT variants attempt to reduce resource requirements, primarily via GPU mem-

ory consumption and storage requirements. By applying prefix-tuning and our variation, prefix-propagation to long-document models we limit carbon emissions and increase accessibility for low-resource groups. We note that prefix-propagation neither exacerbates nor alleviates other ethical risks such as biases regarding gender, race, religion, etc. that are often embedded in pre-trained LLMs. If such biases exist in the pre-trained model, they will be propagated to downstream tasks regardless of tuning method.

Acknowledgements

This research is supported by NSERC Discovery Grants. The second author is also supported by the Vector Scholarship in Artificial Intelligence.

References

- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv:2004.05150*.
- Elad Ben Zaken, Yoav Goldberg, and Shauli Ravfogel. 2022. BitFit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1–9, Dublin, Ireland. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Yifan Chen, Devamanyu Hazarika, Mahdi Namazifar, Yang Liu, Di Jin, and Dilek Hakkani-Tur. 2022. Inducer-tuning: Connecting prefix-tuning and adapter-tuning. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. Generating long sequences with sparse transformers.
- Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, David Belanger, Lucy Colwell, and Adrian Weller. 2020. Rethinking attention with performers.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jacob R Gardner, Geoff Pleiss, David Bindel, Kilian Q Weinberger, and Andrew Gordon Wilson. 2018. Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. In *Advances in Neural Information Processing Systems*.
- Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. 2022. PPT: Pre-trained prompt tuning for few-shot learning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8410–8423, Dublin, Ireland. Association for Computational Linguistics.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17*, page 1321–1330. JMLR.org.
- Jun He, Liqun Wang, Liu Liu, Jiao Feng, and Hao Wu. 2019. Long document classification from local word glimpses via recurrent attention learning. *IEEE Access*, 7:40707–40718.
- Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. 2021a. Towards a unified view of parameter-efficient transfer learning.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.
- Xuehai He, Zhuo Cai, Wenlan Wei, Yichen Zhang, Luntian Mou, Eric Xing, and Pengtao Xie. 2021b. Towards visual question answering on pathology images. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 708–718, Online. Association for Computational Linguistics.
- Edward Hu, Yelong Shen, Phil Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models.
- Matthew Jagielski, Om Thakkar, Florian Tramèr, Daphne Ippolito, Katherine Lee, Nicholas Carlini, Eric Wallace, Shuang Song, Abhradeep Thakurta, Nicolas Papernot, and Chiyuan Zhang. 2022. Measuring forgetting of memorized training examples.

- A. Katharopoulos, A. Vyas, N. Pappas, and F. Fleuret. 2020. Transformers are rnns: Fast autoregressive transformers with linear attention. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Johannes Kiesel, Maria Mestre, Rishabh Shukla, Emmanuel Vincent, Payam Adineh, David Corney, Benno Stein, and Martin Potthast. 2019. [SemEval-2019 task 4: Hyperpartisan news detection](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 829–839, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. 2020. [Reformer: The efficient transformer](#).
- Ken Lang. 1995. Newsweeder: Learning to filter netnews. In *Proceedings of the Twelfth International Conference on Machine Learning*, pages 331–339.
- Quentin Lhoest, Albert Villanova del Moral, Patrick von Platen, Thomas Wolf, Mario Šaško, Yacine Jernite, Abhishek Thakur, Lewis Tunstall, Suraj Patil, Mariama Drame, Julien Chaumond, Julien Plu, Joe Davison, Simon Brandeis, Victor Sanh, Teven Le Scao, Kevin Canwen Xu, Nicolas Patry, Steven Liu, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Nathan Raw, Sylvain Lesage, Anton Lozhkov, Matthew Carrigan, Théo Matussière, Leandro von Werra, Lysandre Debut, Stas Bekman, and Clément Delangue. 2021. [Datasets: A Community Library for Natural Language Processing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184. Association for Computational Linguistics.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597, Online. Association for Computational Linguistics.
- Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. 2022. [P-tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 61–68, Dublin, Ireland. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *ArXiv*, abs/1907.11692.
- Mahdi Pakdaman Naeni, Gregory Cooper, and Milos Hauskrecht. 2015. [Obtaining well calibrated probabilities using bayesian binning](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 29(1).
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [PyTorch: An Imperative Style, High-Performance Deep Learning Library](#). In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Hao Peng, Nikolaos Pappas, Dani Yogatama, Roy Schwartz, Noah A. Smith, and Lingpeng Kong. 2021. [Random feature attention](#).
- Uri Shaham, Elad Segal, Maor Ivgi, Avia Efrat, Ori Yoran, Adi Haviv, Ankit Gupta, Wenhan Xiong, Mor Geva, Jonathan Berant, and Omer Levy. 2022. [Scrolls: Standardized comparison over long language sequences](#).
- Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. [Energy and policy considerations for deep learning in NLP](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3645–3650, Florence, Italy. Association for Computational Linguistics.
- Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. 2022. [Efficient transformers: A survey](#). *ACM Comput. Surv.*, 55(6).
- Yao-Hung Hubert Tsai, Shaojie Bai, Makoto Yamada, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019. Transformer dissection: An unified understanding for transformer’s attention via the lens of kernel. In *EMNLP*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2018. [Constructing datasets for multi-hop reading comprehension across documents](#). *Transactions of the Association for Computational Linguistics*, 6:287–302.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [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. Association for Computational Linguistics.

Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. 2020. Big bird: Transformers for longer sequences. *Advances in Neural Information Processing Systems*, 33.

A Kernel Decomposition Derivation

In the unified framework of He et al. (2021b), we can write the first layer $l = 1$ attention mechanism of prefix-propagation as:

$$H_{l,i} = \text{Attn}(\text{cat}(P^{(l)}, C)W_q^{(l)(i)}, \text{cat}(P^{(l)}, C)W_k^{(l)(i)}, \text{cat}(P^{(l)}, C)W_v^{(l)(i)}) \quad (4)$$

where P is a trained prefix for each downstream task. Omitting layer and head indices and using $D = \text{cat}(P, C)$ for brevity, we can rewrite Equation 4 as:

$$\begin{aligned} H &= \text{Attn}(DW_q, \text{cat}(P, C)W_k, \text{cat}(P, C)W_v) \\ &= \text{softmax}(DW_q \text{cat}(PW_k, CW_k)) \begin{bmatrix} PW_v \\ CW_v \end{bmatrix} \\ &= (1 - \lambda(C)) \text{softmax}(DW_q W_k^\top C^\top) CW_v \\ &\quad + \lambda(C) \text{softmax}(DW_q W_k^\top P^\top) PW_v \\ &= (1 - \lambda(C)) \text{Attn}(DW_q, CW_k, CW_v) \\ &\quad + \lambda(C) \text{Attn}(DW_q, PW_k, PW_v) \\ &= (1 - \lambda(C)) \text{Attn}(\text{cat}(P, C)W_q, CW_k, CW_v) \\ &\quad + \lambda(C) \text{Attn}(\text{cat}(P, C)W_q, PW_k, PW_v) \end{aligned} \quad (5)$$

where $\lambda(C)$ is a scalar (dependent on C) to normalize softmax over the sequence and the prefixes and is computed by:

$$\lambda(C) = \frac{\sum_i DW_q W_k^\top P^\top}{\sum_i DW_q W_k^\top P^\top + \sum_j DW_q W_k^\top C^\top} \quad (6)$$

We consider the two terms of Equation 5 as kernelized attention modules which brings us back to the complete kernel decomposition:

$$H = \text{Kern}(\text{cat}(P, C)W_q, CW_k, CW_v) + (\alpha) \text{Kern}(\text{cat}(P, C)W_q, PW_k, PW_v) \quad (7)$$

where α is an introduced hyperparameter that replaces the fixed weighting of λ . This change allows us to explicitly increase the weighting of prefixes

Artifact	Version	License
transformers (Wolf et al., 2020) ³	4.23.1	Apache 2.0
datasets (Lhoest et al., 2021) ⁴	2.6.1	Apache 2.0
GPyTorch (Gardner et al., 2018) ⁵	1.9.0	MIT
RoBERTa (Liu et al., 2019) ⁶	base	MIT
Longformer (Beltagy et al., 2020) ⁷	base	Apache 2.0
P-Tuning (Liu et al., 2022) ⁸	2.0	Apache 2.0
ArXiv (He et al., 2019) ⁹	no_ref	Unspecified
Hyperpartisan (Kiesel et al., 2019) ¹⁰	1.0	CC BY 4.0
20-newsgroup (Lang, 1995) ¹¹	1.0	Unspecified
WikiHop (Welbl et al., 2018) ¹²	1.1	CC BY SA 3.0

Table 4: Complete list of artifacts used in our experiments along with their versions and licenses.

by scaling the prefix kernel’s coefficients. Kern is the kernelized attention variant described in Tsai et al. (2019):

$$\text{Kern}(Q, K, V)_i = \sum_{j=1}^N \frac{k(Q_i, K_j)}{\sum_{j'=1}^N k(Q_i, K_{j'})} V_j \quad (8)$$

where subscripts (e.g. i) index the rows of a matrix, N is the number of key and value vectors, and k is a kernel function that calculates the similarity score between two vectors. We do not experiment with altering the kernel type since the default exponential kernel inherent to softmax attention already implicitly maps the input vectors to an infinite feature space. Therefore, the kernel function in Equation 8 takes the form:

$$k(x_q, x_k) = \exp\left(\frac{\langle x_q, x_k \rangle}{\sqrt{d_k}}\right) \quad (9)$$

where $\langle \cdot, \cdot \rangle$ signifies the dot product and d_k is the dimension of key projections.

B Experimental Details

Artifact Notes Table 4 summarizes the complete list of artifacts we used in our experiments along with their licenses and versions. All libraries were used for their intended purpose of open-source development. The ArXiv, Hyperpartisan, and WikiHop datasets were released in research contexts to

³<https://github.com/huggingface/transformers>

⁴<https://github.com/huggingface/datasets>

⁵<https://github.com/cornellius-gp/gpytorch>

⁶https://huggingface.co/docs/transformers/model_doc/roberta

⁷https://huggingface.co/docs/transformers/model_doc/longformer

⁸<https://github.com/THUDM/P-tuning-v2>

⁹<https://huggingface.co/datasets/ccdv/arxiv-classification>

¹⁰<https://github.com/zliucr/hyperpartisan-news-detection>

¹¹https://scikit-learn.org/1.2/datasets/real_world.html

¹²<http://qangaroo.cs.ucl.ac.uk/>

	HY.	ArXiv	NG.	WikiHop
Epochs	20	10	10	10
Prefix Len	8	8	8	32
Batch Size			32	
Dropout			0.1	
LR Warmup			0.1	
LR Schedule			linear	
Vocab size			50265	
Loss			cross-entropy	
Optimizer			AdamW	

Table 5: Hyperparameters used for our experiments. “HY.” and “NG.” denote the Hyperpartisan task and the 20-newsgroups tasks, respectively.




Figure 3: Violin plot showing distribution of sequence lengths for each dataset. “HY.” and “NG.” denote the Hyperpartisan task and the 20-newsgroups tasks, respectively.

evaluate and/or develop state-of-the-art algorithms. The intended use of 20-newsgroups is not explicit, although it is commonly used for natural language processing in research. We therefore believe we have adhered to the intended usages of the datasets we included.

We do not anonymize the data for 20-newsgroups as (a) the trained models is not being deployed (only used for evaluation purposes) and (b) the non-anonymized variant is already publicly available. We chose to use the datasets in the current form for fair comparison with other baselines and therefore did not do a detailed analysis for those artifacts. We refer readers to the cited original works in Table 4 for complete documentation.

Training For our experiments, we use and adapt the prefix-tuning implementation provided in Liu et al. (2022). Training was conducted on 12 NVIDIA GeForce 1080 Ti cards, for an estimated 2300 single GPU hours (including preliminary ex-

periments). All models tested fit on a single card, so we did not use any model parallelism. Throughout experiments, we use gradient accumulation for an effective batch size of 32. We use early stopping for our hyperparameter search, and show results for the run with the best validation F1-score. For learning rate, we search between $\{1e-2, 5e-2, 1e-3, 5e-3, 5e-4\}$ for prefix-based methods, and $\{3e-5, 5e-5\}$ for fine-tuning. For kernelized prefix-propagation, we search for a scale factor (hyperparameter α) of $\{1e-2, 4e-2, 1e-3, 3e-3, 5e-3, 7e-3\}$ (after choosing the best learning-rate). Other hyperparameters are listed in Table 5.

Despite seeding random number generators for Hugging Face’s transformer library through the `set_seed` method, slight deviations will propagate if using GPUs due to some non-deterministic CUDA methods that do not respect the seed setting mechanisms of Pytorch (Paszke et al., 2019). Upon further analysis, we found this issue in non-deterministic algorithms to be widely overlooked in the field, and believe that this area needs further discussion in the research community. However, we note that our results should be reproducible when running across multiple seeds.

Task Details All datasets used have a considerable portion of documents greater than RoBERTa’s max sequence limit of 512 tokens, as shown in Figure 3. Number of samples and number of classes for each dataset are in Table 6.

For all classification tasks, we prepend a globally-attended [CLS] token to the start of the sequence and pass the output into a learned classification head. We truncate document lengths to 4096 and 512 tokens for Longformer and RoBERTa, respectively. For Hyperpartisan, we use the same data pre-processing and training split as Beltagy et al. (2020). However, we noticed overlap between training and testing samples, so we instead show validation results. We use the ArXiv dataset from He et al. (2019) that is available on Huggingface datasets (which we reviewed for correctness). The original dataset has labels leaked in the source text, so we use the `no_ref` version that has those labels filtered. We use the 20-newsgroups and follow pre-processing as recommended by scikit-learn authors, removing headers, quotations, and signatures from each sample to prevent the model from learning spurious correlations.

WikiHop instances include a question, candidate answers, and multiple context documents. For

Dataset	n_{sample}	n_{class}	$n_{train/dev/test}$
HY.	645	2	80/10/10
NG.	18,846	20	60/20/20
ArXiv	33,388	11	85/7.5/7.5
WikiHop	48,867	—	90/5/5

Table 6: Datasets used and their total size (n_{sample}), number of classes (n_{class}), and relative sizes of train, validation, and test splits ($n_{train/dev/test}$).




Figure 4: Calibration (measured by ECE) of different tuning approaches using Longformer-base on ArXiv. Lower is better.

a fair comparison, we follow the WikiHop setup in Beltagy et al. (2020) to the best of our ability. In summary, we pass the dataset fields into the model in the format: [q] <question> [/q] [ent] <candidate 1> [/ent] ... [ent] <candidate N> [/ent] [sep] <context 1> [sep] ... [sep] <context N>. Because the context documents are often longer than the maximum sequence length of Longformer, we split the context documents into chunks of 4096 (or 512 for RoBERTa) and pass them separately through the model while concatenated to the question and candidate pair. We then train a classifier to predict a single logit for each [ent] token, take the average over all chunks, apply softmax, and finally use cross-entropy loss. We also train the new special tokens [ent] and [q] in prefix-based methods to better learn an effective representation (as they did not appear in pre-training).

C Impact of Training Time on ECE

Apparent in Figure 4, prefix-propagation is better-calibrated relative to other approaches throughout training. Prefix-tuning and fine-tuning however

Method	Absolute Runtime (s)	Relative Runtime
No PEFT	2192	0%
Prefix-Tuning	2239	+2.1%
Prefix-Propagation	2196	+0.2%

Table 7: Runtime for inference using “No PEFT” (i.e., regular forward pass), prefix-tuning, and prefix-propagation. “Relative Runtime” is the runtime relative to “No PEFT”.

either start less calibrated or deviate from prefix-propagation as training progresses.

D Runtime Performance

We test the inference time of the studied methods and show the results in Table 7. We use the same 8000 randomly generated sequences of length 4096 across methods and test on a NVIDIA GTX 1080 Ti. We notice that prefix-propagation is slightly more efficient than prefix-tuning. We theorize that this discrepancy is caused by prefix-propagation only needing to concatenate a matrix in the first layer (and sum on the rest), whereas prefix-tuning concatenates before every layers.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
See the limitations section on page 5.
- A2. Did you discuss any potential risks of your work?
See the ethics statement on page 5.
- A3. Do the abstract and introduction summarize the paper’s main claims?
See the abstract and introduction (Section 1).
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Yes, we develop our own code for our experiments (Section 4.1). The code will be released.

- B1. Did you cite the creators of artifacts you used?
We provided a complete list of artifacts along with citations in Appendix B (Table 4).
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
See Appendix B (Table 4).
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
See Appendix B.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
See Appendix B.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
We did not create dataset artifacts, so we did not feel it was necessary. We wanted to use the datasets as-is to get a sense of how well our variation performs compared to other models. Doing an in-depth analysis of the dataset was beyond the scope of our paper. See Appendix B for an in-paper explanation.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
See Appendix B (Table 6).

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

C Did you run computational experiments?

See section 4.

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?

See section 4.1 and Appendix B.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

See Appendix B.

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

In various places, for example, in Table 1, Figure 2, and Limitations.

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

See Appendix B.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.