

A Single Linear Layer Yields Task-Adapted Low-Rank Matrices

Hwichan Kim^{†*}, Shota Sasaki[‡], Sho Hoshino[‡], Ukyo Honda[‡]

[†]Tokyo Metropolitan University, [‡]CyberAgent

^{‡‡} Tokyo, Japan

kim-hwichan@ed.tmu.ac.jp, {sasaki_shota, hoshino_sho, honda_ukyo}@cyberagent.co.jp

Abstract

Low-Rank Adaptation (LoRA) is a widely used Parameter-Efficient Fine-Tuning (PEFT) method that updates an initial weight matrix W_0 with a delta matrix ΔW consisted by two low-rank matrices A and B . A previous study suggested that there is correlation between W_0 and ΔW . In this study, we aim to delve deeper into relationships between W_0 and low-rank matrices A and B to further comprehend the behavior of LoRA. In particular, we analyze a conversion matrix that transform W_0 into low-rank matrices, which encapsulates information about the relationships. Our analysis reveals that the conversion matrices are similar across each layer. Inspired by these findings, we hypothesize that a single linear layer, which takes each layer's W_0 as input, can yield task-adapted low-rank matrices. To confirm this hypothesis, we devise a method named Conditionally Parameterized LoRA (CondLoRA) that updates initial weight matrices with low-rank matrices derived from a single linear layer. Our empirical results show that CondLoRA maintains a performance on par with LoRA, despite the fact that the trainable parameters of CondLoRA are fewer than those of LoRA. Therefore, we conclude that "a single linear layer yields task-adapted low-rank matrices." The code used in our experiments is available at <https://github.com/CyberAgentAILab/CondLoRA>.

Keywords: Pretrained Language Model, Parameter-Efficient Fine-Tuning, Low-Rank Adaptation

1. Introduction

In natural language processing (NLP) area, it is common practice to fine-tune pre-trained language models (PLMs) (Devlin et al., 2019; Lewis et al., 2020; Brown et al., 2020) using task-specific data. As the scale of these PLMs has grown considerably, the computational resources required for fine-tuning all parameters have escalated, presenting a substantial computational challenge. In recent years, parameter-efficient fine-tuning (PEFT) methods, which use a limited number of additional parameters, have been proposed to address this issue. PEFT methods include prompt-tuning (Lester et al., 2021), prefix-tuning (Li and Liang, 2021), and low-rank adaptation (LoRA) (Hu et al., 2022), etc. These methods reduce computational costs to fine-tune PLMs while achieving comparable performance to fine-tuning all of the parameters.

Among the PEFT methods, LoRA has been prominent in NLP area because it shows stable and good performance across various NLP tasks and PLMs (Pu et al., 2023). LoRA fixes an initial weight matrix W_0 and updates W_0 with a delta matrix ΔW consisting of trainable low-rank matrices A and B , significantly reducing the number of trainable parameters compared to fine-tuning all parameters. Subsequent studies (Zhang et al., 2023; Valipour et al., 2023) have analyzed several aspects of LoRA to achieve potentially more effective and efficient PLM fine-tuning. Hu et al. (2022)

performed an analysis of the relationship between W_0 and trained $\Delta W (= BA)$, and they revealed that there is a correlation between W_0 and ΔW . This finding implies the existence of certain relationships between the initial weight matrix W_0 and the low-rank matrices A and B .

In this study, we conduct an in-depth analysis of the relationships between initial weight matrices W_0 and low-rank matrices A and B to gain a deeper understanding of LoRA behavior. Specifically, we analyze a conversion matrix that transforms W_0 into A or B under the assumption that it roughly represents their relationships. Our analysis shows that similarities between each layer's conversion matrix are very high. This empirical observation implies a commonality in the relationships between the initial weight matrices and low-rank matrices regardless of layers. Inspired by the results, we hypothesize that a single linear layer, which takes each layer's W_0 as input, can produce task-adapted low-rank matrices of each layers.

To confirm our hypothesis, we design a method named Conditionally Parameterized LoRA (CondLoRA) that fine-tune PLMs with low-rank matrices derived from a single linear layer (Figure 2). Our experiments demonstrate that CondLoRA achieves competitive performance compared to LoRA in GLUE tasks. Notably, CondLoRA can reduce the number of trainable parameters compared to LoRA, because its parameters are constant regardless of target layers. The success of CondLoRA suggests potential avenues for further minimization of trainable parameters in LoRA variants. Our contributions in this study are twofold:

*This work was carried out during his CyberAgent internship period.

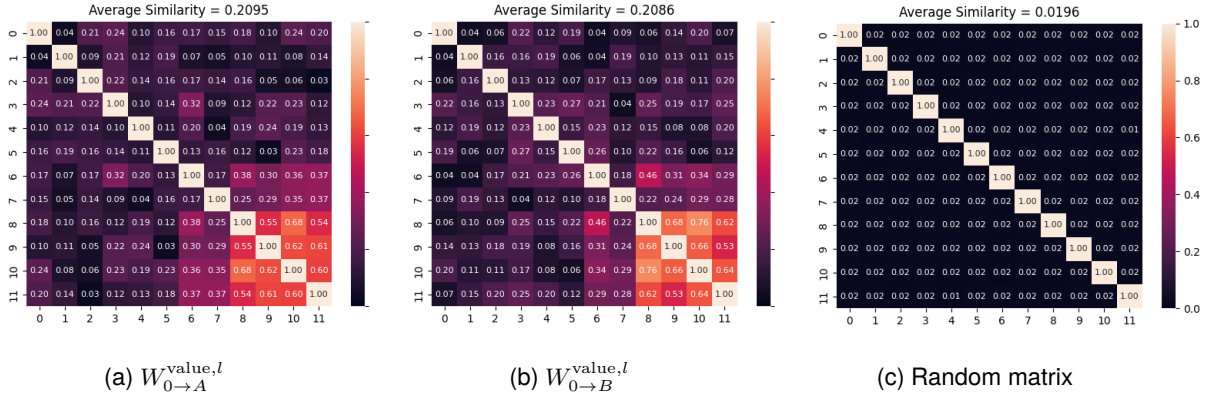


Figure 1: Normalized subspace similarities between each layer’s conversion matrices and random matrices. Average similarity refers to the average of elements excluding the diagonal elements.

Hyperparameters	Value
Batch Size	16
Optimizer	Adam
Scheduler	Linear
Target Modules	{query, value}
Target Layers	{1, 2, ..., 12}
LoRA r	8
LoRA α	8
Max Seq. Len.	512

Table 1: Hyperparameters used in our experiments.

1. We reveal that conversion matrices that transform initial weight matrices into trained low-rank matrices are similar across each layer, which means that there is similar relationship regardless of layers.
2. We demonstrate that CondLoRA achieves performance comparable to the already parameter-efficient LoRA with fewer parameters. This outcome suggests that task-adapted low-rank matrices can be yielded by a single linear operation.¹

2. Preliminaries for LoRA

A diff-planing method (Guo et al., 2021) updates an initial weight matrix $W_0^{m, l} \in \mathbb{R}^{d_1 \times d_2}$ using an trainable matrix $\Delta W^{m, l} \in \mathbb{R}^{d_1 \times d_2}$. Where $m \in \{m_1, \dots, m_k\}$ and $l \in \{1, 2, \dots, N\}$ are target module (e.g., query, value, etc.) and layer, respectively, and N is a total number of layers. Hu et al. (2022) proposed a PEFT method named LoRA, which decompose $\Delta W^{m, l}$ into two low-rank weight matrices:

$$W_0^{m, l} + \Delta W^{m, l} = W_0^{m, l} + B^{m, l} A^{m, l} \quad (1)$$

¹We will publish the code used in our experiments.

where $A^{m, l} \in \mathbb{R}^{r \times d_1}$ and $B^{m, l} \in \mathbb{R}^{d_2 \times r}$ with $r \ll d_2, d_1$. $A^{m, l}$ and $B^{m, l}$ are trained by downstream-task data. Their experiments demonstrated that LoRA achieves comparable or even better performance than full fine-tuning while reducing the numbers of trainable parameters.

In addition, they analyzed several aspects of trained $A^{m, l}$, $B^{m, l}$, and $\Delta W^{m, l}$ using normalized subspace similarity. They defined normalized subspace similarity $\phi(\cdot)$ between matrices $X \in \mathbb{R}^{d_1^X \times d_2^X}$ and $Y \in \mathbb{R}^{d_1^Y \times d_2^Y}$ as:

$$\phi(X, Y, i, j) = \frac{\|U_X^{i\top} U_Y^j\|_F^2}{\min(i, j)} \in [0, 1] \quad (2)$$

where U_X is a left or right unitary matrix and U_X^i is top- i singular vectors of U_X . Therefore, when $\phi(X, Y, i, j)$ is close to 1, it means the singular vector spaces between X and Y are similar. Hu et al. (2022) measured subspace similarities between each $W_0^{m, l}$ and $\Delta W^{m, l}$, and showed that the similarities are higher than those of random Gaussian matrices. This result suggests there are relationships between an initial weight matrix $W_0^{m, l}$ and low-rank matrices $A^{m, l}$ and $B^{m, l}$.

3. Commonality of Relationships across Layers

In this study, we conduct an in-depth analysis of the relationships between initial weight matrices W_0 and low-rank matrices A and B to comprehend LoRA behavior. To analyze the relationships, we use a conversion matrix that transform $W_0^{m, l}$ to $A^{m, l}$ or $B^{m, l}$, under the assumption that it roughly represents their relationships. The analyses of the conversion matrix are expected to provide a deeper understanding of the relationship.

Conversion matrices $W_{0 \to A}^{m, l}$ and $W_{0 \to B}^{m, l}$ satisfy $W_0^{m, l} W_{0 \to A}^{m, l} = (A^{m, l})^\top$ and $W_0^{m, l} W_{0 \to B}^{m, l} = B^{m, l}$,

	MNLI	SST-2	MRPC	CoLA	QNLI	QQP	RTE	STS-B	Avg.
LoRA	86.6	93.7	86.2	61.2	92.0	90.5	74.3	89.3	83.38
CondLoRA	86.5	93.8	86.6	61.1	91.8	90.1	74.2	90.3	83.42
Δ	-0.1	0.1	0.4	-0.1	-0.2	-0.4	-0.1	1.0	0.04

Table 2: Evaluation on GLUE tasks. These scores are the average of three models.

respectively.² $(A^{m,l})^\top$ is a transposed matrix of $A^{m,l}$. Therefore, the conversion matrices are:

$$W_{0 \rightarrow A}^{m,l} = (W_0^{m,l})^{-1} (A^{m,l})^\top \in \mathbb{R}^{d_2 \times r} \quad (3)$$

$$W_{0 \rightarrow B}^{m,l} = (W_0^{m,l})^{-1} B^{m,l} \in \mathbb{R}^{d_2 \times r} \quad (4)$$

where $(W_0^{m,l})^{-1}$ is an inverse matrix of $W_0^{m,l}$.

In this study, we investigate the similarity of conversion matrices across layers. Specifically, we measure normalized subspace similarities (Equation 2) between conversion matrices of each layer. When the similarities are high, it suggests there is a similar relationship between initial weight matrices $W_0^{m,l}$ and low-rank matrices $A^{m,l}$ and $B^{m,l}$ across layers.

3.1. Experimental Settings

We used RoBERTa base (Liu et al., 2019) as a base model and HuggingFace Transformers (Wolf et al., 2020).³ We used PEFT library⁴ and a single NVIDIA A100 40GB for LoRA tuning. We fine-tuned the model using GLUE (Wang et al., 2018) dataset. We set the hyperparameters except for learning rates following Hu et al. (2022) as shown in Table 1. We fine-tuned the models using only 90% of training set. We allocated the remaining 10% for development and used the official GLUE development set as our test data because GLUE dataset does not provide an official test set. We set max epochs to 50 in MNLI and QQP and 100 in other tasks, respectively. Based on evaluation scores in development data, we searched learning rates through Optuna (Akiba et al., 2019)⁵ and selected the best checkpoint. For evaluation metrics, we used the overall (matched and mismatched) accuracy for MNLI, Matthew’s correlation (Matthews, 1975) for CoLA, Pearson correlation for STS-B, and accuracy for other tasks. To measure the normalized subspace similarity (Equation 2), we used a left unitary matrix and set i and j to be 8 ($= r$).

²Notably, we assume that d_1 and d_2 are the same value in our experiments, because we used RoBERTa base and target modules are query and value. If d_1 and d_2 are different values, it is necessary to apply some extensions for calculating a conversion matrix.

³<https://huggingface.co/roberta-base>

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

⁵<https://optuna.org/>




Figure 2: An overview of CondLoRA.

3.2. Experimental Results

Figure 1 shows normalized subspace similarities between conversion matrices (Equations 3 and 4) of each layer and those of random Gaussian matrices. Due to the limited number of pages, we only show the similarities of conversion matrices from a model fine-tuned by MNLI dataset and value modules. The similarities of conversion matrices were higher than those of random matrices. This result implies a commonality in the relationships between the initial weight matrices $W_0^{m,l}$ and low-rank matrices $A^{m,l}$ and $B^{m,l}$ regardless of layers. Inspired by this result, we hypothesize that a single linear layer, which takes each layer’s $W_0^{m,l}$ as input, can produce task-adapted low-rank matrices $A^{m,l}$ (or $B^{m,l}$) of each layer.

In addition, this analysis reveals another noteworthy observation that the similarities between the deeper layers are extremely high. We would like to investigate the underlying causes in future work (See Section 5 for details).

4. Can a Single Linear Layer Yield Task-Adapted Low-Rank Matrices?

In this section, to confirm our hypothesis (Section 3), we design a method named Conditionally Parameterized LoRA (CondLoRA) that fine-tune PLMs

X	Y	l -th layer											
		1	2	3	4	5	6	7	8	9	10	11	12
$A^{\text{value},l}$	$A_{\text{cond}}^{\text{value},l}$	0.05	0.05	0.06	0.04	0.04	0.08	0.03	0.04	0.03	0.06	0.05	0.07
$B^{\text{value},l}$	$B_{\text{cond}}^{\text{value},l}$	0.04	0.03	0.03	0.05	0.05	0.04	0.04	0.03	0.02	0.07	0.08	0.10
$\Delta W^{\text{value},l}$	$\Delta W_{\text{cond}}^{\text{value},l}$	0.04	0.03	0.03	0.05	0.05	0.05	0.04	0.03	0.02	0.07	0.09	0.12

Table 3: Normalized subspace similarity between matrices from LoRA and CondLoRA.

	Trainable parameters	Speed (examples/s)
LoRA	294,912	39.652
CondLoRA	24,576	40.303

Table 4: Trainable parameters and speed during training.

with low-rank matrices derived from a single linear layer. CondLoRA finds low-rank matrices $A_{\text{cond}}^{m,l}$ and $B_{\text{cond}}^{m,l}$ using linear layers as follows:

$$A_{\text{cond}}^{m,l} = \text{Linear}(W_0^{m,l}; \theta_A^m)^\top \in \mathbb{R}^{r \times d_1} \quad (5)$$

$$B_{\text{cond}}^{m,l} = \text{Linear}((W_0^{m,l})^\top; \theta_B^m) \in \mathbb{R}^{d_2 \times r} \quad (6)$$

$$\Delta W_{\text{cond}}^{m,l} = B_{\text{cond}}^{m,l} A_{\text{cond}}^{m,l} \quad (7)$$

where $\theta_A^m \in \mathbb{R}^{d_2 \times r}$ and $\theta_B^m \in \mathbb{R}^{d_1 \times r}$ are trainable parameters. CondLoRA train θ_A^m and θ_B^m using downstream task data. We provide an overview of CondLoRA as shown in Figure 2.

One of the advantages of CondLoRA is its ability to decrease the numbers of trainable parameters. LoRA requires $(d_1 \times r + d_2 \times r) \times k \times N$ trainable parameters. However, CondLoRA requires $(d_1 \times r + d_2 \times r) \times k$ trainable parameters regardless of N , because it use a linear layer per target modules and low-rank matrices. To substantiate our hypothesis, we conduct a comparative analysis of LoRA and CondLoRA based on their performance in GLUE tasks.

4.1. Experimental Results

Table 2 shows the evaluation scores of validation data in each task.⁶ The average scores (Avg.) across all the tasks are nearly equal between LoRA and CondLoRA. Furthermore, CondLoRA outperforms LoRA in SST-2, MRPC, and STS-B tasks. We also performed a pairwise t -Test to measure the statistical significance of the performance difference. The p -values were over 0.01 in all the

⁶We used the same settings in Subsection 3.1.

tasks, indicating that CondLoRA achieves competitive performance compared to LoRA. From the experimental results, we conclude that “a single linear layer yields task-adapted low-rank matrices”.

4.2. Analysis

The numbers of trainable parameters. As explained at Section 4, the numbers of trainable parameters of CondLoRA is constant regardless of the number of target layers. We show the numbers of trainable parameters of CondLoRA and LoRA in Table 4. Table 4 shows that CondLoRA reduces the numbers of trainable parameters to $\frac{1}{12}$ compared to LoRA, because RoBERTa base is consisted by 12 layers and we targeted all layers.

Speed. CondLoRA has extra calculations compared to LoRA, because it determines $A_{\text{cond}}^{m,l}$ and $B_{\text{cond}}^{m,l}$ based on $W_0^{m,l}$ (Equations 5 and 6). There is no difference in inference speed between LoRA and CondLoRA, since the calculations are performed only once when loading a model. However, during training, CondLoRA may takes extra time compared to LoRA because the calculations are required per each iteration. Therefore, we quantified speeds, the numbers of examples processed per second, during training of both LoRA and CondLoRA as shown in Table 4. Contrary to expectations, CondLoRA is slightly faster than LoRA. We consider that the delay from the calculations for low-rank matrices are offset by the backpropagation process, because the trainable parameters (i.e. the parameters to be updated by backpropagation) are fewer than LoRA.

Similarity between low-rank matrices. Finally, we measured normalized subspace similarity between $A^{\text{value},l}$, $B^{\text{value},l}$ and $A_{\text{cond}}^{\text{value},l}$, $B_{\text{cond}}^{\text{value},l}$, respectively. We used right and left unitary matrices for A and B , respectively, and set 8 as i and j . Table 3 demonstrates that while the similarities are not exceedingly high, they are higher than those of random Gaussian matrices.⁷ A similar trend was also observed in query modules. This result implies

⁷The similarities between random matrices are less than 0.01.

that LoRA and CondLoRA, to some degree, obtain similar low-rank matrices.

5. Conclusion and Future Work

In this study, we demonstrated that similar relationships exist between the initial weight matrices and low-rank matrices regardless of layers. Inspired by this analysis, we designed CondLoRA, which updates initial weights with low-rank matrices derived from a single linear layer. We showed that CondLoRA achieves competitive performance compared to LoRA, while trainable parameters are reduced.

Although out of scope of this study, we believe that CondLoRA has the potential to reduce the trainable parameters of other LoRA variants such as AdaLoRA (Zhang et al., 2023). Therefore, for future work, we would like to apply CondLoRA to other LoRA variants and investigate its effectiveness (See Section 5 for details).

Limitations

Although our analyses provided novel insights to achieve more effective and efficient PLMs fine-tuning, our current work has the following limitations.

Analyses in other models and tasks. We used RoBERTa base and GLUE tasks to conduct the analyses of Sections 3 and 4. It is not immediately clear whether these conclusions would hold true for other PLMs and tasks. Therefore, for future work, we are interested in analyzing other PLMs and various tasks to verify if similar results can be achieved irrespective of PLMs or tasks.

Analyses of conversion matrices In this study, we have investigated whether conversion matrices are similar across each layer. We also observed that the similarities between the deeper layers are extremely high. Phang et al. (2021) fine-tuned all of the PLM’s parameters with task-specific data and measured similarities between each layer. They showed that the similarities between the deeper layers are high compared to others. We would like to investigate that a causal relationship between layer similarities of the model fine-tuned all of the parameters and those of conversion matrices. In addition, we have not conducted an analysis of the conversion matrix itself. More detailed analyses about the conversion matrix will provide further insight into the nature of these relationships.

Additionally, Equations 3 and 4 are not able to use to matrices where d_1 and d_2 are different, as inverse matrices are not able to find for such matrices. Therefore, in order to conduct analyses using conversion matrices for other modules, such as

feed-forward layers, we aim to devise a more generalized method for finding conversion matrices in future work.

Evaluation of CondLoRA with LoRA variants

CondLoRA, a method that finds low-rank matrices for each layer using a linear layer, can also be applied to other LoRA variants. For instance, CondLoRA could be applied to AdaLoRA (Zhang et al., 2023), decomposing $\Delta W^{m,l}$ into the form of singular value decomposition $P^{m,l}\Lambda^{m,l}Q^{m,l}$, using trainable parameters θ_P^m , θ_Λ^m , and θ_Q^m for each matrix. Therefore, further investigation into the effectiveness of CondLoRA when applied to other LoRA variants remains a challenge for future research.

Ethical Consideration

In recent years, bias in data has become an issue. Training a models, not just on training using CondLoRA, on biased datasets can result in unjustified predictions or the generation of pejorative content directed towards specific individuals or groups. Intuitively, if the parameters are abundant, the effect of bias will be distributed across each parameter, but if they are small, all parameters may be affected by bias. Since CondLoRA has very small trainable parameters compared to other fine-tuning methods (such as full parameter tuning and the other PEFT methods), it may be more susceptible to the effects of bias. Therefore, when using CondLoRA, sufficient attention should be paid to the problem of bias.

Acknowledgments

We would like to thank Toshio Hirasawa for his help with proposing CondLoRA and writing our paper.

6. Bibliographical References

- Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- 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.
- 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.
- Demi Guo, Alexander Rush, and Yoon Kim. 2021. [Parameter-efficient transfer learning with diff pruning](#). 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 4884–4896, Online. Association for Computational Linguistics.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. 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.
- 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 pre-training approach](#). *CoRR*, abs/1907.11692.
- B.W. Matthews. 1975. [Comparison of the predicted and observed secondary structure of t4 phage lysozyme](#). *Biochimica et Biophysica Acta (BBA) - Protein Structure*, 405(2):442–451.
- Jason Phang, Haokun Liu, and Samuel R. Bowman. 2021. [Fine-tuned transformers show clusters of similar representations across layers](#). In *Proceedings of the Fourth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 529–538, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- George Pu, Anirudh Jain, Jihan Yin, and Russell Kaplan. 2023. [Empirical analysis of the strengths and weaknesses of PEFT techniques for LLMs](#). In *ICLR 2023 Workshop on Mathematical and Empirical Understanding of Foundation Models*.
- Mojtaba Valipour, Mehdi Rezagholizadeh, Ivan Kobzyev, and Ali Ghodsi. 2023. [DyLoRA: Parameter-efficient tuning of pre-trained models using dynamic search-free low-rank adaptation](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3274–3287, Dubrovnik, Croatia. Association for Computational Linguistics.
- Qingru Zhang, Minshuo Chen, Alexander Bukharin, Pengcheng He, Yu Cheng, Weizhu Chen, and Tuo Zhao. 2023. [Adaptive budget allocation for parameter-efficient fine-tuning](#). In *The Eleventh International Conference on Learning Representations*.

7. Language Resource References

- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi 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 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.