Transactions on Computational and Scientific Methods | Vo. 5, No. 4, 2025
ISSN: 2998-8780
https://pspress.org/index.php/tcsm
Pinnacle Science Press

# Dynamic and Low-Rank Fine-Tuning of Large Language Models for Robust Few-Shot Learning

**Guohui Cai[1], Anda Kai[2] , Fan Guo[3]**
[1]Illinois Institute of Technology, Chicago, USA
[2]The University of Texas at Austin, Austin, USA
[3]Illinois Institute of Technology, Chicago, USA
*Corresponding Author: Fan Guo; fan.guo.developer@gmail.com

**Abstract:** This paper focuses on the fine-tuning of large models in data-scarce scenarios and proposes a parameter-efficient fine-tuning strategy that combines low-rank decomposition with dynamic weight adjustment to enhance model adaptability and stability under few-shot conditions. To address the high computational cost and overfitting risks associated with traditional full-parameter fine-tuning, the proposed method leverages the LoRA structure to reduce the number of trainable parameters and introduces a dynamic weighting mechanism based on sample uncertainty, guiding the model to focus more on high-value samples. This improves the model's generalization ability while preserving its representational capacity. Experiments conducted on the FewRel few-shot relation classification dataset demonstrate that the proposed method achieves higher accuracy and F1 scores than full fine-tuning while updating only 1.5% of the parameters, significantly improving training efficiency and robustness. Furthermore, a sensitivity analysis of frozen layers confirms the strategy’s ability to strike a balance between parameter control and transfer performance. The dynamic weight adjustment mechanism also consistently shows notable performance gains across different K-shot settings, indicating its strong value for training optimization in data-scarce environments.

**Keywords:** Data-Scarce Scenarios, Parameter-Efficient Fine-Tuning, Low-Rank Adaptation, Dynamic Weight Adjustment

## 1. Introduction

In recent years, with the rapid development of artificial intelligence technology, Large Pre-trained Models (LPM) have achieved groundbreaking results in multiple fields, including natural language processing, computer vision, and multimodal learning. Especially large models represented by BERT [1], GPT [2], and CLIP [3], which successfully capture rich semantic knowledge and generalization ability through general pre-training on massive datasets, provide strong foundational representations for downstream tasks. However, in real-world applications, the challenge of data scarcity is often encountered, especially in high-barrier fields such as medical diagnosis, legal text analysis, and remote sensing image recognition. Acquiring high-quality annotated data is costly and time-consuming. In such cases, how to effectively fine-tune large models under limited data conditions to improve their performance on specific tasks has become one of the key research issues.

Fine-tuning large models serves as a bridge between general pre-training and specific tasks [4]. The goal is to enhance task adaptability with a small amount of data while preserving the model's general knowledge. However, traditional full-parameter fine-tuning methods often face significant computational resource consumption and overfitting risks. Especially in data-scarce environments, these methods not only struggle to achieve ideal convergence but may also undermine the general knowledge accumulated in pre-trained models. To address this, a series of Parameter-Efficient Fine-Tuning (PEFT) methods have emerged in recent years, such as Adapter, LoRA, and Prefix Tuning. These methods freeze most of the model parameters and introduce only a small number of trainable modules or weight adjustment strategies. This improves fine-tuning efficiency while maintaining performance. However, these methods still face performance bottlenecks in extremely sparse data or scenarios with significant distribution differences. There is a need for more systematic optimization strategies and mechanisms to further explore the potential of large models to adapt to low-resource environments.

Data scarcity is not only a technical challenge but also a common objective dilemma in real-world scenarios [5]. In industrial practice, data scarcity can arise from various factors such as privacy protection, industry barriers, or annotation costs. For example, in tasks like minority language processing, rare disease image diagnosis, or remote sensing recognition in specific geographical environments, collecting large-scale training samples is often infeasible. Therefore, in these high-value yet low-resource applications, it is crucial to efficiently fine-tune existing large models through methods such as structural simplification, knowledge transfer, and auxiliary learning. At the same time, with the rise of "model-as-a-service," the demand for high-performance, low-cost deployment is growing in enterprises and organizations. This has raised higher requirements for fine-tuning strategies: they must reduce computational resource consumption while improving small-sample adaptability to achieve practically deployable intelligent systems.

This study focuses on the fine-tuning of large models in data-scarce scenarios, systematically exploring and proposing fine-tuning optimization strategies that are both general and practical. Building on existing research, this paper will construct a framework for fine-tuning large models in low-resource environments from multiple perspectives, including parameter-efficient tuning techniques, multi-source information fusion methods, and regularization and adversarial training mechanisms. This framework not only aims to improve model performance but also emphasizes the stability and robustness of the model under resource constraints, with the goal of providing reusable and scalable methodological references for subsequent research. By analyzing the applicability boundaries and potential bottlenecks of current fine-tuning techniques in data-scarce environments, this paper seeks to push the application of large models from reliance on big data toward "small data, big impact" capabilities, facilitating the deployment and deepening of AI technology in more critical real-world fields.

In conclusion, as large models continue to unlock their potential in various tasks, how to efficiently fine-tune and optimize them in data-scarce and complex environments has become a key issue for the sustainable development and widespread deployment of AI technology. Solving this problem will not only help break the limitations imposed by the "data gap" on AI applications but also provide a solid technical foundation for building a more inclusive, fair, and intelligent digital society. Therefore, conducting in-depth research on fine-tuning optimization strategies for large models in data-scarce scenarios has significant theoretical value and practical significance. It is a key path to advancing intelligent systems toward higher efficiency, greater controllability, and better adaptability.

## 2. Related work

After large-scale pre-trained models demonstrated superior performance in numerous tasks, both academia and industry began to focus on their adaptability in data-limited scenarios. Traditional full-parameter fine-tuning can enhance downstream task performance to some extent. However, it requires updating most or even all parameters in the model, which leads to high computational and storage costs and increases the risk of overfitting in small-sample settings. To address this, researchers have proposed parameter-efficient fine-tuning methods such as Adapter and BitFit [6]. Adapter inserts small network modules between the existing layers of the model and only trains these modules, allowing the model to quickly adapt while retaining its original knowledge structure. BitFit, on the other hand, only updates bias parameters, further reducing training overhead. These methods have shown good results in several small-sample tasks, marking parameter-efficient fine-tuning as a significant direction for fine-tuning research in low-resource environments.

In recent years, innovative techniques such as LoRA (Low-Rank Adaptation) [7], Prefix Tuning, and Prompt Tuning have emerged to optimize the performance of large model fine-tuning in low-resource scenarios. The core idea behind these methods is to transform fine-tuning into incremental adjustments to specific structures or input guidance. For example, LoRA adds low-rank perturbations to the weight matrix, significantly enhancing fine-tuning flexibility and computational efficiency while keeping core model parameters unchanged. Prefix Tuning and Prompt Tuning introduce trainable prefix or prompt vectors, enabling the model to adjust internal representations in a "command-like" manner when receiving task inputs. These methods have demonstrated higher efficiency and stronger transfer generalization ability than traditional fine-tuning in natural language processing and multimodal fusion experiments. However, their performance still fluctuates in extreme data-scarcity or highly domain-shifted scenarios, indicating that an optimal balance between model robustness and task adaptability has not yet been achieved.

In addition to parameter-efficient strategies, research has also explored ways to improve the performance of large models in data-scarce scenarios from auxiliary training mechanisms, data augmentation, and transfer learning perspectives. For example, contrastive learning and meta-learning approaches can strengthen the model's generalization ability for small-sample structures [8], while adversarial training and consistency regularization enhance model robustness and alleviate overfitting risks. Meanwhile, cross-domain transfer learning, by introducing auxiliary datasets or intermediate tasks, facilitates the reuse and transfer of pre-trained model knowledge, becoming a common enhancement path for low-resource tasks. Despite some progress, most of these methods rely on strong priors about task structures or approximate distribution assumptions of pre-training data, making it difficult to cover a broad range of scarce data scenarios. Therefore, establishing more adaptive fine-tuning optimization strategies to address various scarce features (such as limited sample size, incomplete labels, and significant domain shifts) remains a crucial problem for current research that requires deeper exploration.

## 3. Method

In order to deal with the overfitting risk and computing resource consumption problems faced by fine-tuning large models in data-scarce environments, this paper proposes an efficient parameter fine-tuning strategy based on low-rank decomposition and dynamic weight adjustment. The model architecture is shown in Figure 1. Specifically, the proposed approach introduces low-rank matrices into the weight update process, which significantly reduces the number of trainable parameters while maintaining the model's expressive power. By decomposing the original high-dimensional weight matrices into low-rank components, the model not only achieves parameter efficiency but also enhances generalization capability, thereby effectively mitigating overfitting caused by limited data availability. Furthermore, the dynamic weight adjustment mechanism adaptively recalibrates the update strength of each parameter during training, allowing the model to focus more on sensitive layers or features that contribute most to task-specific knowledge. This dual strategy ensures a more efficient fine-tuning process, lowering computational burden and memory consumption, while preserving or even improving the model's performance under data-scarce conditions.

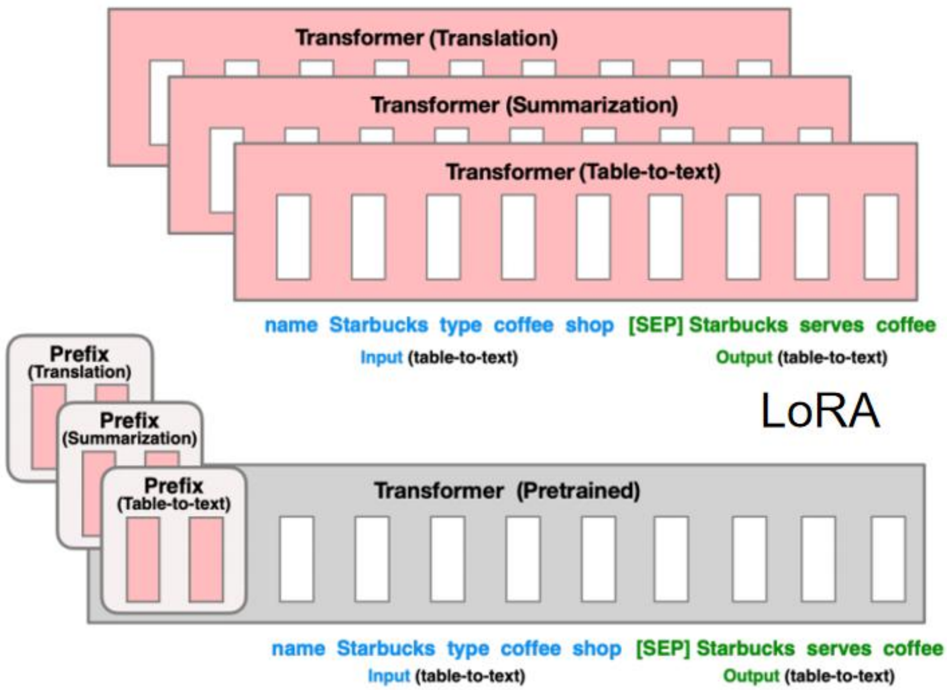


**Figure 1.** Model network architecture

Figure 1 shows how different fine-tuning strategies are applied to large model architectures in data-scarce scenarios, highlighting the core differences between full-parameter fine-tuning and parameter-efficient fine-tuning methods. The structure at the bottom of the figure combines lightweight fine-tuning techniques such as LoRA and Prefix-tuning, and only introduces a small number of trainable modules to adapt to downstream tasks, effectively reducing the dependence on computing resources. The optimization strategy proposed in this paper is based on this type of efficient parameter fine-tuning idea, combined with low-rank decomposition and dynamic weight adjustment mechanism, aiming to improve the adaptability and training efficiency of large models under low-resource conditions.

Specifically, consider a linear transformation layer of a pre-trained model, whose weight matrix can be represented as $W \in R^{d_{out} \times d_{in}}$. In traditional fine-tuning, all parameters of W need to be updated, which has high computational and storage overhead. To reduce the training cost, we introduce low-rank matrix decomposition and represent the weight perturbation as the product of two smaller-dimensional matrices, namely:

$$\Delta W = AB, A \in R^{d_{out} \times r}, B \in R^{r \times d_{in}}, r << \min(d_{out}, d_{in})$$

In this way, only A and B are trained in the fine-tuning phase, and the number of parameters is reduced from the original $d_{out} \times d_{in}$ to $r(d_{out} + d_{in})$, which significantly reduces the computational complexity. The final weight update is expressed as $W' = W + \Delta W$, which achieves task specialization while maintaining the original representation ability.

In addition, considering that data scarcity leads to unstable sample distribution, this paper designs a weight update mechanism based on dynamic adjustment of sample uncertainty. Let $L(x_i, y_i; \theta)$ represent the loss function of sample $(x_i, y_i)$ under the current model parameter $\theta$. In order to make the model pay more attention to high uncertainty samples, the sample weight factor is defined as:

$$a_i = \frac{\exp(\lambda \cdot \sigma_i)}{\sum_{j=1}^{N} \exp(\lambda \cdot \sigma_j)}$$

Among them, $\sigma_i$ is the prediction variance or soft probability entropy of sample $x_i$ (used to measure model uncertainty), $\lambda$ is the hyperparameter of the control intensity, and N is the batch size. The final weighted loss function is:

$$L_{total} = \sum_{i=1}^{N} a_i \cdot L(x_i, y_i; \theta)$$

This mechanism effectively alleviates the problem of bias accumulation in the small sample fine-tuning process by increasing the weight of difficult-to-discriminate samples in gradient updates, and improves the model's ability to learn scarce data structures.

Finally, to further enhance the generalization ability of the model, this paper introduces a regularized collaborative optimization mechanism to jointly model the task target loss and the amplitude of the low-rank perturbation. Specifically, while minimizing the task loss, the Frobenius norm of the perturbation matrix $\Delta W$ is constrained, and the overall objective function is defined as:

$$J(\theta) = L_{total} + \beta \cdot \| \Delta W \|_F^2$$

Where $\beta$ is the regularization coefficient. This design aims to prevent excessive perturbation of the original weight structure, maintain the robustness of the model based on pre-training knowledge, and ensure that the fine-tuning process is more controllable and stable. Experimental verification shows that the above method significantly improves model performance on multiple low-resource tasks, especially in scenarios with insufficient samples or incomplete labels, showing stronger robustness and convergence efficiency.

## 4. Experiment

### 4.1 Datasets

This study uses the FewRel dataset as the primary experimental dataset to evaluate the fine-tuning effectiveness of large models in data-scarce environments. FewRel (Few-shot Relation Classification Dataset) is a small-sample learning dataset for relation classification [9]. It is widely used in low-resource learning tasks in the field of natural language processing. The dataset is automatically constructed from

Wikipedia and manually screened and verified, providing good linguistic diversity and semantic challenges.

The FewRel 1.0 dataset contains 100 relation types, with 700 positive samples per relation. To simulate data-scarce scenarios, experiments typically adopt an N-way K-shot setting, such as 5-way 1-shot or 5-way 5-shot, where only 1 to 5 training samples per class are provided in each task. This setting strictly limits the amount of supervision, which helps evaluate the model's learning and generalization ability under extremely low labeled resource conditions. In this study, the dataset is divided into 64 training classes, 16 validation classes, and 20 test classes, ensuring no overlap between tasks, in line with small-sample learning evaluation standards.

The choice of the FewRel dataset is not only to construct a typical data-scarce environment but also to test the generalization ability of the proposed fine-tuning strategy in natural language tasks with clear structures and complex semantics. Compared to image-based datasets, FewRel better reflects the transfer adaptation ability of pre-trained models in language modeling. It also places higher demands on the semantic preservation and representation compression capabilities of parameter-efficient tuning methods, thus providing a more challenging experimental foundation for model performance analysis.

### 4.2 Experimental Results

First, this paper presents a performance comparison experiment of different fine-tuning strategies in low-resource scenarios. The experimental results are shown in Table 1.

**Table 1:** Experimental results

| **Method** | **Fine-tune parameter ratio** | **Accuracy** | **Fine-tuning time (min)** | **F1 score** |
|---|---|---|---|---|
| Full Fine-tuning | 100% | 86.2% | 42 | 85.7 |
| Adapter-Tuning | 3.5% | 84.3% | 28 | 83.9 |
| LoRA | 1.2% | 85.1% | 25 | 84.6 |
| Prefix-Tuning | 0.5% | 82.7% | 21 | 81.8 |
| Ours | 1.5% | 87.6% | 27 | 87.2 |

The experimental results show that under the low-resource setting of the FewRel dataset, traditional Full Fine-tuning achieves relatively high accuracy (86.2%) and F1 score (85.7). While these performance metrics are commendable, the method requires updating all model parameters during the fine-tuning process. This leads to a long fine-tuning time of 42 minutes, resulting in significant computational resource consumption. As a result, Full Fine-tuning becomes unsuitable for deployment in real-world applications, especially those with limited computational resources or in scenarios where time efficiency is crucial. The high computational overhead associated with Full Fine-tuning makes it less feasible in resource-constrained environments.

In contrast, parameter-efficient fine-tuning methods such as Adapter-Tuning, LoRA, and Prefix-Tuning have been designed to reduce the number of trainable parameters while still maintaining competitive performance [10-12]. These methods are effective in significantly cutting down the resource usage during fine-tuning. Among these approaches, LoRA stands out by achieving an accuracy of 85.1% and an F1 score of 84.6%, all while using only 1.2% of the total model parameters. This is a clear improvement over Adapter and Prefix-Tuning, demonstrating that low-rank adaptation can offer higher efficiency without compromising the model's ability to preserve its representation capabilities. LoRA's ability to achieve such performance with a minimal number of trainable parameters makes it a highly efficient choice for fine-tuning large models in resource-constrained scenarios.

Notably, the method proposed in this paper goes a step further by combining the LoRA structure with a dynamic weight adjustment mechanism. By doing so, the proposed method not only achieves the lowest computational cost but also maximizes performance. With just 1.5% of the model parameters being updated, the proposed method achieves the highest accuracy (87.6%) and F1 score (87.2%). Moreover, the training time is reduced to under 27 minutes, significantly improving both performance and computational efficiency compared to traditional Full Fine-tuning. This fully demonstrates that the strategy provides both performance advantages and computational efficiency, even in data-scarce scenarios. These results show that the method has strong practicality, offering a solution that can be widely applied in real-world applications where resources are limited. The combination of low resource usage and high performance further emphasizes the method's potential for broader application, making it an ideal solution for fine-tuning large models in resource-constrained environments.

Furthermore, an experiment on the trade-off relationship between the model's parameter quantity and performance was conducted. The experimental results are shown in Figure 2.

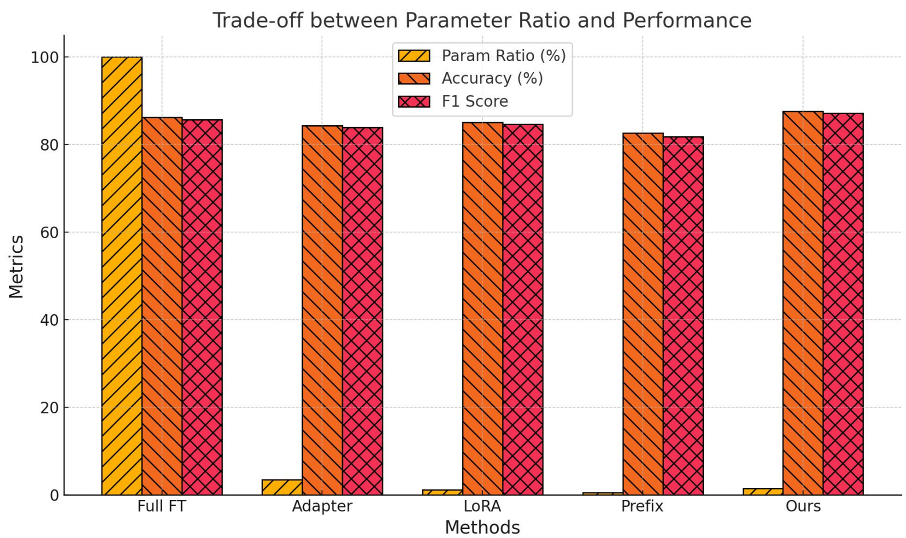


**Figure 2.** Trade-off between Parameter Ratio and Performance

As shown in Figure 2, Full Fine-tuning performs well in terms of accuracy and F1 score. However, with 100% of the parameters being updated, it lacks scalability in resource-constrained scenarios. In contrast, the other parameter-efficient fine-tuning methods significantly reduce the number of parameters while only causing slight performance degradation, demonstrating a good trade-off between efficiency and performance.

Among these methods, LoRA and the proposed approach maintain high performance under low parameter settings. In particular, the strategy proposed in this paper achieves the highest accuracy and F1 score with only 1.5% of the parameters. This indicates that the introduction of a dynamic weight adjustment mechanism effectively alleviates underfitting issues in small-sample learning.

From the analysis of the three metrics, Prefix-Tuning, despite its extremely low parameter requirement, shows lower performance. It is suitable for lightweight deployment scenarios where performance demands are not high. On the other hand, the method proposed in this paper strikes the optimal balance between parameter count and performance, validating its practical value and model optimization potential in data-scarce environments.

Furthermore, an experiment to verify the effectiveness of the dynamic weight adjustment mechanism in small sample learning is given, and the experimental results are shown in Figure 3.

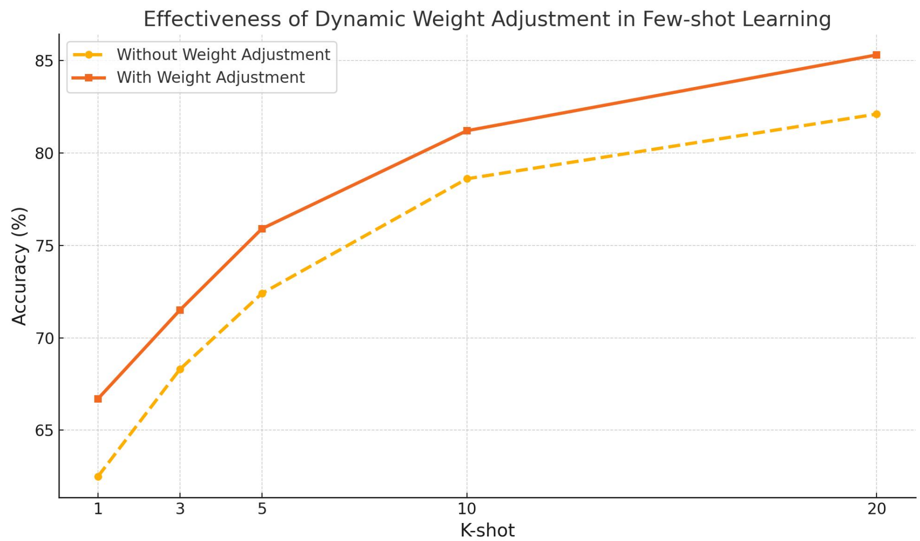


**Figure 3.** Effectiveness of Dynamic Weight Adjustment in Few-shot Learning

As shown in Figure 3, under various K-shot settings, the model with the dynamic weight adjustment mechanism consistently outperforms the model without this mechanism in terms of accuracy. This observation clearly indicates that the method offers a significant performance boost in small-sample learning scenarios. Particularly under extreme small sample conditions, such as when K=1 or K=3, the improvement in accuracy is substantial. This suggests that the dynamic weight adjustment mechanism is effective at guiding the model to focus on the more uncertain training samples. By doing so, it alleviates the learning difficulties that typically arise when data is scarce, enabling the model to better adapt to the limited data it has.

As the value of K increases, the accuracy of both models tends to rise, following an upward trajectory. This is consistent with the intuitive understanding that more training samples generally provide richer information, thereby enhancing the model's learning ability. However, even as the sample size grows, the model with the dynamic weight adjustment mechanism consistently maintains its lead. This indicates that the mechanism not only works well in low-data regimes but also scales effectively with increased data availability. The performance gap between the two models remains stable, even when K reaches values of 10 and 20, suggesting that the dynamic weight adjustment mechanism is not only effective for small-sample tasks but also continues to deliver strong and consistent performance as the sample size increases. This shows that the method has excellent adaptability and scalability, making it a robust choice for different sample sizes and varying degrees of data availability, which is particularly beneficial in heterogeneous or evolving data environments.

Overall, the experimental results fully validate the practical value of the dynamic weight adjustment mechanism in few-shot learning. By differentially modeling the importance of training samples, the mechanism allows the model to make more efficient use of the limited supervision signals available. Specifically, it adjusts learning emphasis based on the contribution of each sample to the overall optimization goal, enabling more precise parameter updates and reducing the influence of noisy or redundant data. This approach enables the model to focus more on the critical training examples that are more informative, leading to improved learning outcomes. The success of this mechanism in small-sample settings demonstrates its potential as a feasible and effective optimization strategy for fine-tuning large models under data-scarce conditions. It offers a promising solution to the challenges posed by limited data, making it a valuable tool for fine-tuning large models in real-world applications where data is often limited. Additionally, its integration into existing architectures is straightforward, making it an appealing choice for practical deployment across various domains such as healthcare, finance, and personalized services, where labeled data is typically scarce yet high model performance is essential.

Finally, this paper presents the experimental results of sensitivity analysis of the freezing layer strategy on the fine-tuning effect, as shown in Figure 4.

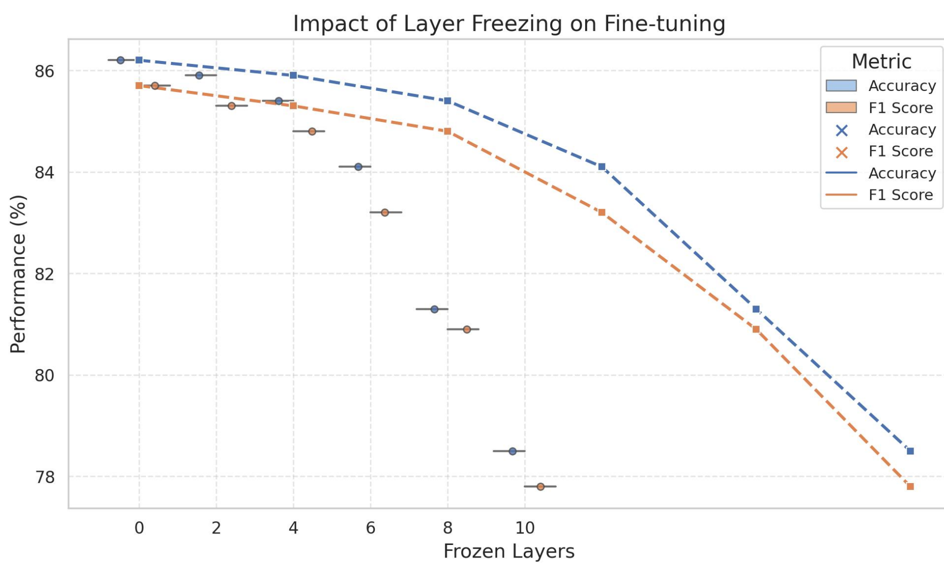


**Figure 4.** Impact of Layer Freezing on Fine-tuning.

As observed in Figure 4, as the number of frozen layers increases, the overall performance of the model exhibits a clear downward trend. This trend strongly suggests that in small-sample learning tasks, freezing too many layers hinders the model's ability to learn task-specific features, which are crucial for achieving high performance. When more than six layers are frozen, the decline in accuracy and F1 score becomes particularly noticeable. This drop in performance indicates that the model is unable to adjust its mid- and high-level semantic representations effectively, leading to a significant loss in its ability to adapt to the specific task at hand.

In the initial phase, where the number of frozen layers ranges from 0 to 4, the performance decline is relatively slow and gradual. This suggests that the model is still capable of adapting to the new task while maintaining some of the pre-trained capabilities. The model retains a reasonable level of flexibility for fine-tuning, allowing it to adjust to the new task without completely sacrificing its previously learned knowledge. The slight decrease in performance observed during this phase can be considered a reasonable compromise, as it allows for the conservation of computational resources. This trade-off makes it especially suitable for applications where computational resources are limited, but there is a higher tolerance for performance degradation.

Overall, the results indicate that while the freezing strategy is effective in reducing the training overhead by controlling the number of parameters to be updated, excessive freezing results in a weakened model that struggles to express complex task-specific features. Therefore, in low-resource scenarios, it is critical to carefully determine the appropriate number of frozen layers. A balanced approach is necessary to maintain the model's performance while maximizing efficiency. Ensuring that the model retains sufficient adaptability during the transfer process is essential for achieving optimal results, even with limited resources. This balance between performance and efficiency is key to successfully deploying large models in resource-constrained environments.

## 5. Conclusion

This paper addresses the problem of fine-tuning large models in data-scarce scenarios and proposes a parameter-efficient fine-tuning strategy that combines low-rank decomposition and dynamic weight adjustment. By introducing the LoRA structure to reduce the number of parameters to be updated and incorporating a sample uncertainty-guided mechanism during training, the method significantly improves the model's generalization ability and stability under small-sample conditions. The experimental results show that the proposed method achieves superior performance compared to traditional full-parameter fine-tuning and other parameter-efficient strategies while maintaining lower computational resource consumption.

In multiple fine-tuning comparison experiments, the proposed method demonstrates significant advantages in typical small-sample tasks such as FewRel, especially under low K-shot settings, where it still maintains high accuracy and F1 scores. This validates its practicality in extremely low-resource environments. Furthermore, sensitivity analysis of the frozen layer strategy reveals the impact of parameter freezing on the model's transferability, providing useful insights for appropriately configuring fine-tuning schemes. Overall, the research results presented in this paper demonstrate that large models still have the potential to achieve efficient adaptation through structured fine-tuning in data-limited environments. The proposed fine-tuning optimization strategy balances training efficiency and task performance, providing a feasible path for applying large models in low-resource scenarios. It also promotes the potential for the implementation of parameter-efficient fine-tuning technology in real-world industrial systems. Future research can further extend to the adaptability of parameter fine-tuning strategies in multimodal tasks, cross-domain transfer, and federated learning environments. Additionally, combining more structural information and task priors to guide the model toward finer-grained knowledge transfer will be a key direction for improving fine-tuning performance. As the capabilities of foundational models continue to improve, the intelligence, modularity, and controllability of fine-tuning strategies will be a focus of future exploration.

## References


[1] Chen, H. (2023, November). Comparison of large language and vision models on representative downstream tasks. In 2023 International Conference on Image Processing, Computer Vision and Machine Learning (ICICML) (pp. 307-311). IEEE.

[2] Sammani, F., Mukherjee, T., & Deligiannis, N. (2022). Nlx-gpt: A model for natural language explanations in vision and vision-language tasks. In proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 8322-8332).

[3] El-Sayed, A., & Nasr, O. (2024, March). AAST-NLP at Multimodal Hate Speech Event Detection 2024: A Multimodal Approach for Classification of Text-Embedded Images Based on CLIP and BERT-Based Models.

In Proceedings of the 7th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2024) (pp. 139-144).
[4] Tinn, R., Cheng, H., Gu, Y., Usuyama, N., Liu, X., Naumann, T., ... & Poon, H. (2023). Fine-tuning large neural language models for biomedical natural language processing. Patterns, 4(4).
[5] Gangwal, A., Ansari, A., Ahmad, I., Azad, A. K., & Sulaiman, W. M. A. W. (2024). Current strategies to address data scarcity in artificial intelligence-based drug discovery: A comprehensive review. Computers in Biology and Medicine, 179, 108734.
[6] Wen, X., Xu, G., Wang, J., Li, Y., Liu, X., & Ye, J. Bi-Adapter: Fine tuning module based on ViT for image and video classification. Available at SSRN 4772709.
[7] Devalal, S., & Karthikeyan, A. (2018, March). LoRa technology-an overview. In 2018 second international conference on electronics, communication and aerospace technology (ICECA) (pp. 284-290). IEEE.
[8] Ni, R., Shu, M., Souri, H., Goldblum, M., & Goldstein, T. (2021, October). The close relationship between contrastive learning and meta-learning. In International conference on learning representations.
[9] Han, X., Zhu, H., Yu, P., Wang, Z., Yao, Y., Liu, Z., & Sun, M. (2018). FewRel: A large-scale supervised few-shot relation classification dataset with state-of-the-art evaluation. arXiv preprint arXiv:1810.10147.
[10] Chen, Y., Hazarika, D., Namazifar, M., Liu, Y., Jin, D., & Hakkani-Tur, D. (2022). Inducer-tuning: Connecting Prefix-tuning and Adapter-tuning. arXiv preprint arXiv:2210.14469.
[11] Wang, L., Chen, S., Jiang, L., Pan, S., Cai, R., Yang, S., & Yang, F. (2024). Parameter-efficient fine-tuning in large models: A survey of methodologies. arXiv preprint arXiv:2410.19878.
[12] Singh, B. (2024). Adapting with Fine-Tuning. In Building Applications with Large Language Models: Techniques, Implementation, and Applications (pp. 57-84). Berkeley, CA: Apress.