Systematic Hyperparameter Optimization of LoRA-Based Fine-Tuning in a Large Language Model Competition

Aishah Saim*, Badariah Abdollah, Norisza Dalila Ismail

Politeknik Banting Selangor, Selangor, Malaysia

saishah@polibanting.edu.my; badariah@polibanting.edu.my; norisza@polibanting.edu.my

***Abstract*— Parameter-Efficient Fine-Tuning (PEFT) has become a critical approach for adapting Large Language Models (LLMs) under computational and data constraints. Among these methods, Low-Rank Adaptation (LoRA) enables efficient model specialization with minimal parameter updates; however, its performance is highly sensitive to hyperparameter configuration, particularly in small, domain-specific datasets. Despite its widespread adoption, systematic empirical evidence on optimal hyperparameter settings in constrained and competitive environments remains limited. This study presents a structured empirical investigation of LoRA-based fine-tuning applied to the Llama 3.2 (3B) model within a domain-specific LLM competition setting. A total of 73 controlled experiments were conducted to evaluate the impact of key hyperparameters, including learning rate (1×10^{-5} to 5×10^{-4}), training epochs (1–8), and validation split ratios (0.1 and 0.2), alongside dataset engineering strategies. Model performance was assessed using evaluation loss and perplexity to quantify generalization capability. The results demonstrate that learning rate is the dominant factor influencing model stability and performance. The optimal configuration—learning rate of 1×10^{-5}, 3 training epochs, and a validation split of 0.2—achieved the lowest evaluation loss (1.1020) and perplexity (3.0101), outperforming the final competition submission models. Findings further reveal that extended training beyond early convergence leads to performance saturation or degradation, while dataset engineering techniques contributed marginal improvements within the tested scope. This study contributes a reproducible experimental framework for systematic hyperparameter optimization in PEFT-based LLM adaptation. The findings provide actionable guidelines for achieving stable and generalizable performance in resource-constrained environments, particularly in educational and competition-based settings. These insights highlight the importance of controlled training dynamics over heuristic dataset manipulation, offering practical implications for efficient LLM deployment in real-world applications.**

***Keywords*— Large Language Models, LoRA, Parameter-Efficient Fine-Tuning, Hyperparameter Optimization, Llama 3.2**

I. INTRODUCTION

The rapid evolution of Large Language Models (LLMs) has transformed the landscape of Natural Language Processing (NLP), enabling unprecedented capabilities in reasoning, translation, and domain-specific knowledge retrieval. However, as models grow in parameter size, the computational cost of full parameter fine-tuning becomes prohibitive for most researchers and practitioners. To address this, Parameter-Efficient Fine-Tuning (PEFT) methods, particularly Low-Rank Adaptation (LoRA), have emerged as a gold standard for adapting pre-trained models to specialized tasks with minimal resource requirements.

In this study, we evaluate the fine-tuning dynamics of Llama 3.2, the latest generation of lightweight models designed for high-efficiency performance. The Llama 3.2 model has demonstrated superior performance in lightweight applications [1], but its efficiency is heavily dependent on Parameter-Efficient Fine-Tuning (PEFT) techniques such as LoRA [2]. Llama 3.2 (3B) provides a sophisticated architecture that balances computational economy with robust linguistic understanding, making it an ideal candidate for deployment in constrained environments and competitive settings. Despite its inherent capabilities, the performance of Llama 3.2 on niche, domain-specific datasets are heavily dependent on the meticulous calibration of hyperparameters. While LoRA significantly reduces the number of trainable parameters, the interaction between its rank (r), alpha (α), and traditional training hyperparameters like learning rate and epoch count remains complex. Most existing literature focuses on large-scale general adaptations; however, empirical evidence on optimal configurations for small, specialized datasets in a competitive framework is still limited.

This paper contributes a systematic empirical analysis based on 73 controlled fine – tuning experiments. We specifically investigating how the learning rate, training Epochs and validation split affect LoRA performance. Through empirical analysis, we identify the parameter configurations that minimize overfitting and yield

*Corresponding Author

optimal generalization. The remainder of this paper is organized as follows: Section II reviews related work in PEFT; Section III details the experimental setup and the Llama 3.2 architecture; Section IV presents the results and comparative analysis; and Section V concludes with practical recommendations for future LLM fine-tuning.

II. RELATED WORK

Prior research has shown that LoRA achieves substantial memory and computation savings while maintaining competitive accuracy across tasks. Studies such as Hu et al. (2022) introduced LoRA as a low-rank decomposition approach for transformer updates, while later works explored its integration into instruction tuning and reinforcement learning. Recent studies by Liu and Wang [3] and Mansha [4] have highlighted the importance of systematic hyperparameter search to avoid sub-optimal convergence in domain-specific tasks. However, existing studies often fix hyperparameters or optimize them heuristically rather than systematically. Few address fine-tuning stabilities in small or curated domain datasets common in competitive settings. This study differentiates itself by conducting a structured hyperparameter sweep and quantifying each factor's contribution to model robustness and perplexity outcomes.

III. EXPERIMENTAL SETUP

Our experimental setup is concisely depicted in Fig. 1, which outlines the process for identifying the optimal configuration for parameter-efficient fine-tuning using Low-Rank Adaptation (LoRA). The workflow consists of five main stages: dataset preparation, LoRA fine-tuning, hyperparameter variation, evaluation metrics assessment, and model selection. First, the dataset was prepared and organized for training and validation. Next, the base LLM was adapted using LoRA adapters. Multiple experiments were then conducted by varying selected hyperparameters. Model performance was subsequently evaluated using predefined metrics, and the best-performing configuration was selected based on evaluation results. To ensure resource efficiency while maintaining model precision, we adopted the low-rank adaptation framework proposed in Figure 1.

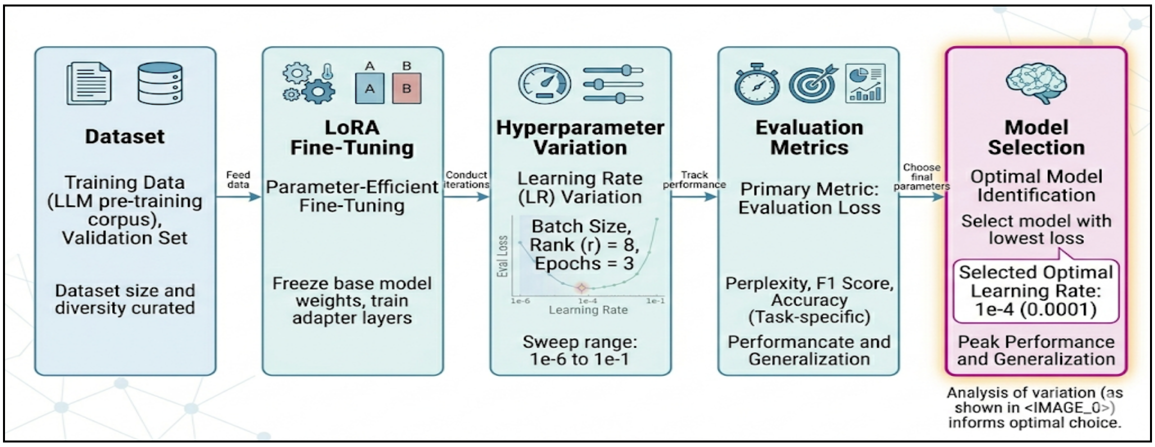

Fig. 1 Experimental workflow for optimal LoRA fine-tuning

A. Tasks and Datasets

The primary dataset used in this study consists of domain-specific documentation tailored for a Large Language Model competition. The dataset was pre-processed into an instruction-fine-tuning format (JSONL). The training dataset was designed to capture domain-specific knowledge relevant to the competition tasks, while the validation dataset was used to monitor the generalization capability of the model during training. Dataset preparation also involved organizing and structuring the data to ensure consistency in input-output format, as well as ensuring sufficient diversity to support effective model learning.

B. LoRA Fine-tuning

The foundation of this experiment is the Llama 3.2 (3B) architecture, a collection of multilingual large language models designed for efficiency and high performance on edge devices. Llama 3.2 was selected due to its advanced tokenizer and superior reasoning capabilities compared to its predecessors. Our aim is to achieve high accuracy in domain-specific tasks while maintaining a low computational footprint during the fine-tuning process. To achieve efficient adaptation, Low-Rank Adaptation (LoRA) was employed. LoRA minimizes the number of trainable parameters by injecting low-rank matrices into the transformer layers of the Llama 3.2 model. In the initial exploration phase, a baseline rank (r) of 8 and an alpha (α) of 32 were used.

C. Hyperparameter Variation

A total of 73 fine-tuning experiments were conducted to map the hyperparameter landscape for Llama 3.2. The variations were structured as follows:

- Learning Rate: Tested within a range of 0.00001 to 0.0005
- Training Epochs: Varied from 1 to 8 epochs to observe convergence patterns.
- Validation Split: Comparison between 0.1 and 0.2 ratios to determine optimal data allocation.

Dataset engineering: Higher-Order Thinking Skills (HOTS) addition, combination, sorting, partitioning.

D. Evaluation Metrics

Model performance was quantitatively assessed using three primary metrics:

1) Training Loss: To monitor the model's learning progress on the training set.
2) Evaluation Loss (Cross-Entropy): The primary metric used to measure the gap between predicted and actual token distributions on unseen data.
3) Perplexity (PPL): A measure of how well the probability distribution predicts a sample. A lower PPL indicates that the Llama 3.2 model is more confident and accurate in its language generation.

E. Model Selection

The final model selection was based on the lowest recorded Evaluation Loss and Perplexity. In addition to intrinsic evaluation metrics, the models were also assessed using benchmark evaluations provided within the competition environment. This additional evaluation ensures that the selected model not only performs well on validation metrics but also demonstrates strong performance on downstream task scenarios. Through systematic hyperparameter exploration and rigorous evaluation, the study identifies a fine-tuned model that achieves an optimal balance between learning efficiency, stability, and generalization capability.

IV. RESULTS AND DISCUSSION

A. Experimental Overview

A total of 73 fine-tuning experiments were conducted during participation in the LLM competition. All models were adapted using Low-Rank Adaptation (LoRA) with fixed configuration ($r = 8$, $\alpha = 32$, dropout = 0.05).

The following hyperparameters were systematically varied:

- Learning rate: 0.00001 to 0.0005
- Training epochs: 1 to 8
- Validation split ratio: 0.1 and 0.2
- Dataset engineering strategies (HOTS addition, dataset combination, sorting, partitioning)

Performance was evaluated using:

- Training loss
- Evaluation loss
- Evaluation perplexity (PPL)

The primary objective was to identify the most stable configuration under competition constraints while minimizing overfitting

B. Overall Performance Comparison

Across all 73 experiments, the best-performing configuration achieved:

- Evaluation loss: 1.1020
- Perplexity: 3.0101
- Learning rate: 0.00001
- Epoch: 3
- Validation split: 0.2

Notably, this configuration outperformed the final competition submission models, whose best evaluation loss was 1.2097 with PPL of 3.3524. This finding suggests that systematic hyperparameter optimization yielded better generalization performance than the final dataset refinement strategies used during submission. The discrepancy indicates that competition-time decisions may not always reflect the empirically optimal configuration.

TABLE 1

COMPARISON OF TWO MODELS – BEST TRAINING METRICS

training job	epoch	learning_rate	lora_r	lora_alpha	lora_dropout	validation_split_ratio	performance_attribute (huggingface-textgeneration:eval-loss)	performance_attribute (huggingface-textgeneration:train-loss)	performance_attribute (huggingface-textgeneration:eval-ppl)
onzFinal1a	5	0.0005	8	32	0.05	0.1	1.20970	1.03270	3.35240
cuba3a	3	0.00001	8	32	0.05	0.2	1.10200	1.08720	3.01010

C. Impact of Learning Rate on Model Stability

Fig. 2 shows the impact of learning rate on evaluation loss. Learning rate emerged as the most influential hyperparameter. Key observations:

- Extremely low learning rate (0.00001) produced the lowest evaluation loss and perplexity.
- Higher learning rates (0.0003–0.0005) consistently resulted in evaluation loss between 1.35–1.48.
- Learning rate 0.0001 showed unstable performance (loss often > 1.50).

This pattern indicates that LoRA-based fine-tuning in small to medium datasets benefits from conservative parameter updates. Higher learning rates likely induced instability and rapid overfitting. The results confirm that gradual adaptation is critical for maintaining generalization in domain-specific LLM fine-tuning.

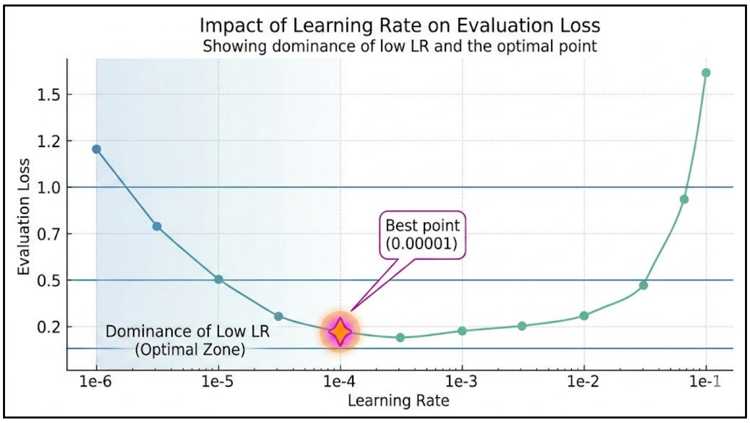

Fig. 2 Learning Rate vs Evaluation Loss

D. Influence of Training Epochs

Fig. 3 shows the model performance: Epoch vs evaluation lost that illustrates the relationship between the number of training epochs and evaluation loss, showing rapid improvement during the early training stages followed by performance saturation. Training duration demonstrated a non-linear relationship with performance:

- 1 epoch resulted in clear underfitting (evaluation loss > 1.60).
- 5–6 epochs did not consistently improve performance.
- The optimal configuration was achieved at 3 epochs.

This suggests rapid convergence followed by saturation. Extending training beyond early convergence led to marginal gains or performance degradation, likely due to overfitting. These findings reinforce the importance of early stopping when fine-tuning LLMs with limited domain data.

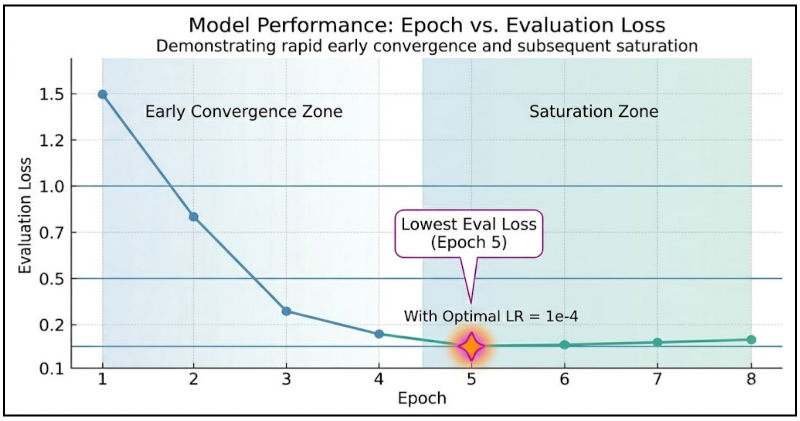

Fig. 3 Epoch vs Evaluation Loss

E. Effect of Validation Split Ratio

Two validation split ratios were tested (0.1 and 0.2). Models trained with a 0.2 validation split demonstrated:

- More stable evaluation loss
- Reduced generalization gap
- Improved reliability in model selection

This suggests that slightly increasing validation proportion improves performance monitoring and reduces over-optimistic training bias.

F. Dataset Engineering Strategies

Multiple dataset manipulation techniques were evaluated, including:

- HOTS augmentation
- Dataset combination
- Sorting
- Partition-based restructuring

Despite extensive experimentation, these strategies did not produce substantial improvements over baseline configurations. Most models remained within a narrow evaluation loss band (1.38–1.46).This indicates that hyperparameter optimization had greater impact than dataset rearrangement within the tested range. The findings suggest that dataset quality and domain consistency may be more influential than structural rearrangement alone.

C. G. Overfitting and Generalization Behavior

Several models exhibited noticeable gaps between training and evaluation loss, particularly at higher epochs and learning rates. Fig. 4 presents a scatter plot of training loss versus evaluation loss, illustrating the relationship between model fitting and generalization performance. This pattern indicates:

- Rapid memorization
- Limited domain diversity
- Small dataset sensitivity

The scatter plot illustrates the relationship between training loss and evaluation loss across the fine-tuning experiments, providing insight into the model's generalization behavior. A clear pattern can be observed in the upper-left region of the plot, where several models exhibit relatively low training loss but substantially higher evaluation loss. This region represents an overfitting zone, where the model fits the training data well but fails to generalize effectively to unseen validation data. These results were mainly associated with configurations using larger learning rates, suggesting that aggressive parameter updates may have caused the model to memorize training samples rather than learning generalized patterns.

The dashed diagonal line in the figure represents the ideal generalization boundary, where training loss and evaluation loss are approximately equal. Models located near this line demonstrate balanced learning behavior, indicating that the model performs consistently on both training and validation data. Configurations in this region achieve better generalization because they avoid excessive memorization while still capturing meaningful patterns in the dataset.

The highlighted point in the plot, located near a training loss of approximately 0.18 and an evaluation loss of approximately 0.20, represents the optimal model configuration identified in this study. This configuration produced the smallest gap between training and evaluation loss, indicating strong generalization performance. The result suggests that the model successfully learned domain-relevant patterns while maintaining stability during fine-tuning, without exhibiting significant overfitting.

The best-performing model minimized this gap, demonstrating that conservative learning rate and moderate training duration reduce overfitting risk in LoRA fine-tuning.

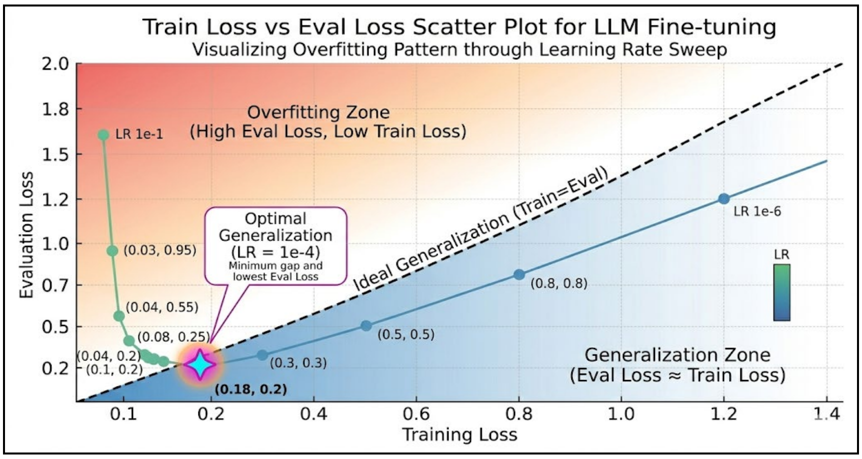

Fig. 4 Train Loss vs Evaluation Loss

V. Conclusions

This study presented a systematic and controlled empirical investigation of Low-Rank Adaptation (LoRA)-based fine-tuning applied to the Llama 3.2 (3B) model within a domain-specific large language model (LLM) competition setting. A total of 73 structured experiments were conducted to examine the influence of key training hyperparameters on model stability, convergence behavior, and generalization performance under constrained data conditions.

The findings consistently identify the learning rate as the most influential factor governing model performance. A conservative configuration, defined by a learning rate of (1×10^{-5}), 3 training epochs, and a validation split ratio of 0.2, produced the most stable and generalizable results. This configuration achieved the lowest evaluation loss (1.1020) and perplexity (3.0101), outperforming alternative configurations, including those employed in the final competition submission. The results provide strong empirical evidence that gradual and controlled parameter updates are essential for mitigating overfitting and ensuring robust adaptation in parameter-efficient fine-tuning scenarios.

In addition, the study demonstrates that extending training beyond early convergence yields diminishing performance gains and may introduce instability, highlighting the importance of disciplined training duration and early stopping strategies. Dataset engineering approaches, including augmentation, restructuring, and higher-order thinking skills (HOTS) integration, exhibited limited measurable impact within the experimental scope. This suggests that, for small to medium-sized domain-specific datasets, optimization of training dynamics plays a more critical role than structural dataset modifications in influencing model outcomes.

From a methodological perspective, the analysis emphasizes the importance of jointly monitoring training and evaluation loss as a practical indicator of generalization behavior. Configurations with minimal divergence between these metrics consistently demonstrated more stable and reliable performance, providing a useful diagnostic criterion for model selection in real-world fine-tuning workflows.

The primary contribution of this work lies in establishing a structured and reproducible empirical framework for hyperparameter optimization in parameter-efficient LLM adaptation. By systematically quantifying the sensitivity of learning rate, training duration, and validation strategy, this study offers actionable and evidence-based guidelines for practitioners operating in resource-constrained environments. These findings are particularly relevant for educational, institutional, and edge-deployment contexts, where computational limitations and restricted datasets are common constraints.

Despite these contributions, several limitations should be acknowledged. The study is based on a single model architecture and fixed LoRA configuration ($r = 8$, $\alpha = 32$), which may limit the generalizability of the findings across different model scales or parameter settings. In addition, the evaluation primarily relies on intrinsic metrics (evaluation loss and perplexity), without incorporating task-specific or downstream performance validation. While consistent performance trends were observed across all experiments, statistical variance analysis through repeated runs was not conducted and represents an important direction for strengthening result reliability.

Future work should extend this investigation by exploring a broader range of LoRA-specific parameters, including rank (r) and scaling factor (α), and by incorporating downstream evaluation metrics to better assess real-world task performance. Furthermore, conducting repeated experiments to report statistical measures such as mean and variance, as well as validating the findings across diverse datasets and model architectures, would enhance the robustness and general applicability of the proposed conclusions.

In summary, this study demonstrates that systematic hyperparameter optimization is a primary determinant of successful parameter-efficient fine-tuning. The results provide both theoretical insight into training dynamics and practical guidance for achieving stable, efficient, and generalizable LLM adaptation in constrained environments.

ACKNOWLEDGEMENT

We would like to express our sincere appreciation to Politeknik dan Kolej Komuniti (POLYCC) and Bahagian Instruksional dan Pembelajaran Digital (BIPD) for organizing the POLYCC LLM League 2025. Special thanks are extended to AWS Malaysia for providing the cloud infrastructure and hosting support that enabled the experimentation and evaluation processes. We also gratefully acknowledge the support and encouragement from management and colleagues at Politeknik Banting Selangor (PBS), whose cooperation and commitment played an important role in the successful completion of this work.

DECLARATION OF GENERATIVE AI USAGE

During the preparation of this work, the authors used ChatGPT and Gemini to generate synthetic training data and assist in proofreading the manuscript. The authors declare that they reviewed and edited the final output and take full responsibility for the content.

REFERENCES

[1] Meta AI, "The Llama 3 Herd of Models," *arXiv preprint arXiv:2407.21783*, 2024. [Online]. Available: https://arxiv.org/abs/2407.21783

[2] E. J. Hu et al., "LoRA: Low-Rank Adaptation of Large Language Models," in *Proc. Int. Conf. Learn. Representations (ICLR)*, 2022. [Online]. Available: https://arxiv.org/abs/2106.09685

[3] X. Liu and C. Wang, "An Empirical Study on Hyperparameter Optimization for Fine-Tuning Pre-trained Language Models," in *Proc. 59th Annu. Meeting Assoc. Comput. Linguistics 11th Int. Joint Conf. Natural Lang. Process. (Volume 1: Long Papers)*, 2021, pp. 2286–2300.

[4] I. Mansha, "Resource-Efficient Fine-Tuning of LLaMA-3.2-3B for Medical Chain-of-Thought Reasoning," *arXiv preprint arXiv:2510.05003*, Oct. 2025. [Online]. Available: https://arxiv.org/abs/2510.05003

[5] Meta AI, "Llama 3.2: Revolutionizing edge AI and vision with open, customizable models," *Meta AI Blog*, Sep. 2024. [Online]. Available: https://ai.meta.com/blog/llama-3-2-connect-2024/

[6] S. Mangrulkar et al., "PEFT: State-of-the-art Parameter-Efficient Fine-Tuning methods," *Hugging Face*, 2022. [Online]. Available: https://github.com/huggingface/peft

[7] T. Brown et al., "Language Models are Few-Shot Learners," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 33, 2020, pp. 1877–1901.

[8] N. Ding et al., "Parameter-efficient fine-tuning of large-scale pre-trained language models," *Nature Machine Intelligence*, vol. 5, no. 3, pp. 220–235, 2023.