

Challenges in Adapting Multilingual LLMs to Low-Resource Languages using LoRA PEFT Tuning

*Omkar Khade^{1,2}, Shruti Jagdale^{1,2}, Abhishek Phaltankar^{1,2}, Gauri Takalikar^{1,2}, Raviraj Joshi^{2,3}

¹Pune Institute of Computer Technology, Pune, India

²Indian Institute of Technology Madras, Chennai, India

³L3Cube Labs Pune, India

Abstract

Large Language Models (LLMs) have demonstrated remarkable multilingual capabilities, yet challenges persist in adapting these models for low-resource languages. In this study, we investigate the effects of Low-Rank Adaptation (LoRA) Parameter-Efficient Fine-Tuning (PEFT) on multilingual Gemma models for Marathi, a language with limited resources. Using a translated Alpaca dataset with 52,000 instruction-response pairs, our findings reveal that while evaluation metrics often show a performance decline post-fine-tuning, manual assessments frequently suggest that the fine-tuned models outperform their original counterparts. The observations indicate improvements in target language generation capabilities but a reduction in reasoning abilities following language adaptation. These results underscore the need for improved evaluation methodologies and the creation of high-quality native datasets to accurately assess language-specific model performance in low-resource settings.

Keywords: LoRA · PEFT · Fine-tuning · Low-resource languages · Marathi · Gemma

1 Introduction

The emergence of Large Language Models (LLMs) such as the Llama and Gemma series has revealed substantial abilities in managing various multilingual tasks (Team et al., 2024a,b). These models have shown competence in multiple high-resource languages, yet their effectiveness with low-resource languages is still a challenge that needs addressing (Huang et al., 2023; Chang et al., 2023). Typically, fine-tuning is used to enhance model performance in particular domains or languages. Nonetheless, this strategy has yielded inconsistent outcomes for low-resource languages (Alam et al., 2024; Lankford et al., 2023a).

Our research focuses on Marathi, which is considered a low-resource language due to the scarcity

of naturally occurring training data (Ogueji et al., 2021; Dhamecha et al., 2021). We leverage the capabilities of LoRA PEFT, a parameter-efficient approach enabling model adaptation, instead of using the classic vanilla Supervised Fine-Tuning (SFT) (Hu et al., 2021; Han et al., 2024). We prefer PEFT over SFT as it works in low data scenarios, is computationally effective so more widely adopted, and avoids catastrophic forgetting due to usage of non-English data only (Weng, 2024; Aggarwal et al., 2024). We execute this method with the Gemma models employing the Alpaca dataset, translated into Marathi. Automated assessments based on NLU and commonsense reasoning usually indicate a decline in the performance of fine-tuned models. However, human evaluations, which directly judge response quality, show that these models excel in specific contextual and cultural aspects (Gala et al., 2024; Zhu et al., 2024).

Our study challenges the effectiveness of current evaluation methods, especially for low-resource languages (Richburg and Carpuat, 2024). We highlight how automated metrics may overlook important qualitative improvements, particularly when models produce responses that resonate with specific linguistic contexts (Barnett et al., 2024). Automated benchmarks, often based on logits, may be unsuitable for evaluating instruction-tuned models, further raising concerns about reliance on these metrics (Gurgurov et al., 2024). We recommend adopting more rigorous evaluation methods that better align with human judgment (Aggarwal et al., 2024; Barnett et al., 2024).

2 Related Work

Using LLMs for low-resource languages, especially Supervised Fine-Tuning (SFT), has been thoroughly researched before. SFT proves to be very effective in high-resource settings, but it falls short in low-resource languages, facing many difficulties

due to the data scarcity. Methods that were curated to handle constraints of low-resource languages were used through multilingual models (Lankford et al., 2023a; Tang et al., 2020). This resulted in highlighting a performance decline, caused by cultural inconsistencies in datasets (Huang et al., 2023; Chang et al., 2023).

As opposed to this, some of the issues have been reduced by parameter-efficient techniques like LoRA PEFT, as they minimize the number of parameters during fine-tuning. This method signifies that computational efficiency is offered and the original model’s robustness is retained, by adjusting only some of the parameters (Hu et al., 2021). A broader study emphasized that using LoRA in low-resource settings comes with low computational overhead (Han et al., 2024; Weng, 2024). Despite this, there remains a considerable gap for exploration when it comes to leveraging LoRA for low-resource languages on Multilingual LLMs (Gurgurov et al., 2024).

Existing frameworks for evaluation of low-resource languages contain limitations that need to be studied (Richburg and Carpuat, 2024; Aggarwal et al., 2024). Low-resource languages have cultural nuances and context-dependent accuracy embedded in them, and traditional evaluation metrics may not capture them (Barnett et al., 2024; Ogueji et al., 2021). This necessitates using alternative evaluation metrics, one of them being human assessments, to corroborate model performance (Gala et al., 2024). For example, as explored, Hindi-language tasks require cultural specificity, as it does for Marathi, our study finds (Dhamecha et al., 2021; Gala et al., 2024). Thus we researched how fine-tuning methods like LoRA produce quality outputs, especially when they are used in culturally refined contexts (Gala et al., 2024; Alam et al., 2024).

3 Experimental Setups

3.1 Dataset

The Alpaca dataset, consisting of 52,000 instruction-response pairs originally in English, was utilized for our research. The Google translate API was used to convert the dataset’s instruction, input, and output columns into Marathi so that it could be used to fine-tune Gemma models. Through this translation process, we were able to produce a sizable dataset for Marathi, which helped us build the models for a language with little resources. The dataset that was created offered a

systematic and uniform format for assessing the performance of the models on instruction-driven tasks in Marathi, making it easier to compare the base and fine-tuned variants of the Gemma models.

3.2 Models and Fine-tuning

For our experiments, we employed several versions of the Gemma model family (Team et al., 2024a) to assess the impact of LoRA PEFT tuning on Marathi, a low-resource language. Specifically, we worked with the following **base models**:

- **gemma-2b**: A 2-billion parameter model with robust multilingual capabilities, serving as one of the baseline models.
- **gemma-2b-it**: An instruction-tuned variant of Gemma-2B, specifically designed to excel at instruction-based tasks.
- **gemma-2-2b**: An enhanced and more recent version with additional pretraining on multilingual corpora, aimed at improving performance in complex linguistic tasks.
- **gemma-2-2b-it**: An instruction-tuned variant of Gemma-2.2B, optimized further for multilingual and instruction-following tasks.

We fine-tuned these models using LoRA PEFT to efficiently adapt them to the Marathi language, producing the following **fine-tuned models**:

- **gemma-2b (Mr)**: The fine-tuned version of Gemma-2b for Marathi using the Alpaca dataset.
- **gemma-2-2b (Mr)**: The fine-tuned version of Gemma-2-2b for Marathi.
- **gemma-2-2b-it (Mr)**: The fine-tuned version of Gemma-2-2b-it for Marathi, specialized for instruction-following tasks.

LoRA PEFT allowed us to tune a smaller subset of model parameters, which minimized computational costs while maintaining the core functionality of the Gemma models. This approach was particularly advantageous in adapting these large models to a low-resource language like Marathi, where we aimed to optimize model performance without requiring extensive computational resources.

MODEL	F1 Scores				
	indicsentiment	ai2_arc-easy	arc challenge	indic copa	indic xnli
gemma-2b	0.7772	0.4435	0.4240	0.6547	0.3582
gemma-2b-it	0.7444	0.4651	0.4043	0.2963	0.3066
gemma-2b (Mr)	0.9397	0.6048	0.3848	0.4219	0.1675

Table 1: F1 Scores for Gemma1 models.

MODEL	F1 Scores				
	indicsentiment	ai2_arc-easy	arc challenge	indic copa	indic xnli
gemma-2-2b	0.9206	0.6384	0.6463	0.6577	0.2191
gemma-2-2b (Mr)	0.8411	0.6135	0.5271	0.5764	0.2753
gemma-2-2b-it	0.9749	0.6851	0.7210	0.7210	0.2814
gemma-2-2b-it (Mr)	0.9589	0.6343	0.6374	0.5835	0.1667

Table 2: F1 Scores for Gemma2 models.

3.3 Evaluation

Our assessment emphasizes two complementary methods:

Automated Evaluation: We utilize established benchmarks from AI4Bharat to assess the performance of the models on tasks such as IndicSentiment, ARC-easy, ARC Challenge, Indic COPA, and Indic XNLI (Gala et al., 2024). These benchmarks enable a quantitative evaluation of the models across a variety of language tasks, allowing us to compare the results with those of other multilingual models

Manual Evaluation: As we used the automated metrics, we also performed thorough assessments manually, using a subset of 150 questions from our curated sheet of questions. Then, leveraging the models, we generated responses for each model and each question to ascertain which model demonstrated better performance. The questions encompassed fields like knowledge-based, quantitative analysis, culture and history, mathematics, science, problem-solving, scenario-based, geography, and politics. This manual evaluation revealed some important model capabilities that were previously overlooked by automated metrics, like cultural significance, linguistic patterns, nuances, and the capacity to follow instructions

By integrating both automated and manual evaluations, we achieved a more thorough understanding of model performance, pinpointing areas where fine-tuned models excel and where they may fall short.

4 Results

4.1 Result Discussion

In the manual assessment of 150 questions, illustrated in Figure 1, fine-tuned versions like gemma-2-2b-it (Mr) and gemma-2b-it (Mr) showed higher win rates than their base counterparts, indicating their enhanced ability to generate contextually relevant answers in Marathi. Nonetheless, the base models occasionally generated responses in English, as depicted in Appendix Figure 2, revealing ongoing issues with language consistency that the fine-tuned models somewhat alleviated, though not completely. While the fine-tuned models performed better in most of the aspects, there were some instances where the base models matched their performance, reflecting the intricate challenges of adapting models for low-resource languages such as Marathi.

In the evaluation of the F1 score, represented in Table 1 for gemma-1 models and Table 2 for gemma-2 models, gemma-2-2b frequently performed better than the other models in significant benchmarks, including sentiment analysis and question-answering tasks. However, fine-tuned models like gemma-2-2b-it (Mr) displayed varied outcomes, showing enhancements in certain tasks while experiencing declines in others, particularly in benchmarks like Indic XNLI and ARC Challenge. These findings highlight that even though fine-tuning can enhance performance in specific areas, it does not guarantee improvements across all tasks, underlining the necessity for more focused fine-tuning strategies for low-resource languages.

Overall, we observe a degradation in NLU and

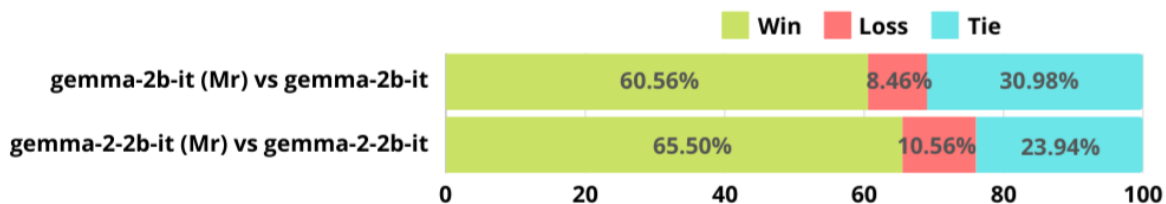


Figure 1: Manual Evaluation Performance.

reasoning benchmarks following language adaptation. However, the adapted model performs better on the open-ended question answering dataset during manual evaluation. This suggests the need for a more comprehensive evaluation strategy and more suitable datasets to fully assess the benefits of language adaptation. While automated benchmarks indicate degradation, they may not be the ideal metric for evaluating instruction-based models. We require more effective benchmarks that can assess the reasoning capabilities of the model without relying on logit-based evaluation metrics.

5 Limitations

While researching, we faced quite a few limitations that hindered progress. Firstly, we used a dataset that was translated, instead of fetching naturally occurring Marathi content from the web. This proved unfruitful as the translated dataset does not entirely capture the complexities of the language. Next issue we faced was of limited computational resources, which resulted in limited experimental explorations, and thwarting us from exploring a broader range of models. Another challenge pertained to comprehensively evaluating the Marathi language generation, as previous benchmarks may not understand its complexities. Furthermore, the translation process contained biases, affecting the accuracy and quality of the question-answer pairs. Lastly, high-quality Marathi evaluation datasets were scarce, limiting our abilities in judging model performance in detail, this called for more robust resources in low-resource settings.

6 Conclusion

To conclude, our results showcase how fine-tuning of Gemma models for Marathi using LoRA PEFT compromises performance if it is based on tradi-

tional and automated evaluation metrics. On the contrary, manual assessments indicate better performance as the fine-tuned models excel in processing culturally sound and contextually relevant responses. This necessitates the use of alternate and enhanced evaluation techniques that can successfully take into account the complex nuances of low-resource languages.

A change needs to be made in developing more robust evaluation methods which provide more accuracy and more effective performance in low-resource settings. Moreover it is also important to perpetuate the generation of high-quality naturally occurring Marathi datasets for continued advancements in this discipline.

7 Acknowledgments

This work was done under the L3Cube Labs, Pune. We would like to express our gratitude towards our mentors at L3Cube for their continuous support and encouragement. This work is a part of the L3Cube-MahaNLP initiative (Joshi, 2022).

References

- Divyanshu Aggarwal, Anuj Sathe, Ian Watts, and Sunayana Sitaram. 2024. [Maple: Multilingual evaluation of parameter efficient finetuning of large language models](#). *ArXiv*.
- Firoj Alam, Shammur Absar Chowdhury, Sabri Boughorbel, and Maram Hasanain. 2024. LLMs for low resource languages in multilingual, multimodal and dialectal settings. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: Tutorial Abstracts*, pages 27–33.
- Samuel Barnett, Zachary Brannelly, Steven Kurniawan, and Samuel Wong. 2024. [Fine-tuning or fine-failing? debunking performance myths in large language models](#). *ArXiv*.

- Tyler A. Chang, Caleb Arnett, Zhezheng Tu, and Benjamin K. Bergen. 2023. [When is multilinguality a curse? language modeling for 250 high- and low-resource languages.](#) *ArXiv*.
- Tejas I. Dhamecha, V. Ramasubramanian Murthy, Smitha Bharadwaj, Karthik Sankaranarayanan, and Pushpak Bhattacharyya. 2021. [Role of language relatedness in multilingual fine-tuning of language models: A case study in indo-aryan languages.](#) *ArXiv*.
- Julian Martin Eisenschlos, Sebastian Ruder, Piotr Czapla, Marcin Kardas, Sylvain Gugger, and Jeremy Howard. 2019. [Multifit: Efficient multi-lingual language model fine-tuning.](#) *ArXiv*.
- Jay Gala, Theja Jayakumar, Javed Ahmad Husain, Arun Kumar M, Mohammed Shakib Khan, Diptesh Kanojia, Ratish Puduppully, Mitesh M. Khapra, Raj Dabre, Radhika Murthy, and Anoop Kunchukuttan. 2024. [Airavata: Introducing hindi instruction-tuned llm.](#) *ArXiv*.
- Daniel Gurgurov, Michael Hartmann, and Simon Ostermann. 2024. [Adapting multilingual llms to low-resource languages with knowledge graphs via adapters.](#) *ArXiv*.
- Zhangyin Han, Cheng Gao, Jiaxin Liu, Jiaqi Zhang, and Sara Qin Zhang. 2024. [Parameter-efficient fine-tuning for large models: A comprehensive survey.](#) *ArXiv*.
- Edward J. Hu, Yelong Shen, Phil Wallis, Zeyuan Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models.](#) *ArXiv*.
- Haoyang Huang, Tianyi Tang, Dongdong Zhang, Wayne Xin Zhao, Tao Song, Yingce Xia, and Furu Wei. 2023. Not all languages are created equal in llms: Improving multilingual capability by cross-lingual-thought prompting. In *Conference on Empirical Methods in Natural Language Processing*.
- Raviraj Joshi. 2022. [L3cube-mahanlp: Marathi natural language processing datasets, models, and library.](#) *ArXiv*.
- S. Lankford, H. Afi, and A. Way. 2023a. [adaptmllm: Fine-tuning multilingual language models on low-resource languages with integrated llm playgrounds.](#) *Information*, 14(12):638.
- S. Lankford, H. Afi, and A. Way. 2023b. [adaptmllm: Fine-tuning multilingual language models on low-resource languages with integrated llm playgrounds.](#) *Information*, 14(12):638.
- Kelechi Ogueji, Yuxin Zhu, and Jimmy Lin. 2021. Small data? no problem! exploring the viability of pretrained multilingual language models for low-resourced languages. In *Proceedings of the 1st Workshop on Multilingual Representation Learning*, pages 116–126.
- Alana Richburg and Marine Carpuat. 2024. [How multilingual are large language models fine-tuned for translation?](#) *ArXiv*.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. [Multilingual translation with extensible multilingual pretraining and finetuning.](#) *ArXiv*.
- Gemma Team, Thomas Mesnard, Cooper Hardin, et al. 2024a. [Gemma: Open models based on gemini research and technology.](#) *ArXiv*.
- Gemma Team, Morgane Riviere, Shivang Pathak, et al. 2024b. [Gemma 2: Improving open language models at a practical size.](#) *ArXiv*.
- Brian Weng. 2024. [Navigating the landscape of large language models: A comprehensive review and analysis of paradigms and fine-tuning strategies.](#) *ArXiv*.
- Dexin Zhu, Pinzhen Chen, Meng Zhang, Barry Haddow, Xianqiang Shen, and Dietrich Klakow. 2024. [Fine-tuning large language models to translate: Will a touch of noisy data in misaligned languages suffice?](#) *ArXiv*.

Appendix

Example Outputs

Question	Model			
	gemma-2b-it	gemma-2b(Mr)	gemma-2-2b-it	gemma-2-2b-it(Mr)
Translate the following sentence to Marathi. India is a great country.	India ँ एक महत्त्वपूर्ण देश।	भारताची एक महान देश आहे.	भारता - means "India" एक - means "a" महत्त्वपूर्ण - means "great" or "important" देश - means "country" अरे - is a common way to express "is" in Marathi.	भारता एक महत्त्वाची देश आहे.
८५ चे १०% किती आहे?	८५ चे १०% किती आहे या मजकरी म्हणता की 85% वरीयता या 10% वरीयता या 0.10 या 10/100 या 0.1 या 10 अक्षरोंचा मिश्रित मजकरी आहे जो 85 या 10 या 0.1 या 10 अक्षरोंचा मिश्रित मजकरी आहे	85 चे 10% 8.5 आहे.	८५ चे १०% आहे ८.५ Calculation: 10% of 85 = (10/100) * 85 = 8.5	८५ चे १०% आहे ८.५.
भारताची सर्वाधिक लांबट नदी कोणती आहे?	हाँ, भारतची सर्वाधिक लांबट नदी कोणती आहे.	भारताची सर्वाधिक लांब नदी नालंदा आहे, जी 2,525 किमी (1565 मैल) लांब आहे.	भारताची सर्वाधिक लांबट नदी नर्मदा आहे. Explanation: The question asks for the longest river in India. The answer is the Narmada River.	भारताची सर्वाधिक लांबट नदी गंगा आहे.

Figure 2: Responses