The 8th International Conference on Interdisciplinary Research
on Computer Science, Psychology, and Education (2024)

Exploring Language Transfer Techniques in Large Language Models

Jungseob Lee[1], Kinam Park[2*]
[1]Department of Computer Science and Engineering, Korea University, [2]Human-inspired AI Research
{omanma1928, spknn}@korea.ac.kr

Abstract

This paper explores various techniques for transferring pre-trained large language models, such as Llama-2, to cross-lingual models. We focus on methods that efficiently enhance performance in languages other than English using minimal data and computational resources. We compare different approaches, including Supervised Fine-tuning (SFT), Direct Preference Optimization (DPO), and Odds Ratio Preference Optimization (ORPO), and discuss their merits and limitations in the context of cross-lingual transfer learning.

Keywords: Large Language Models, Language Adaptation

1. Introduction

Recent advancements in open large language models (LLMs) like Llama [1, 2, 3] have significantly enhanced the ability to understand and generate language similar to humans. These models have demonstrated remarkable capabilities in a wide range of natural language processing (NLP) tasks, including text generation, translation, summarization, and question answering. As a result, there has been considerable interest in extending the utility of these English-centric LLMs to other languages for various applications [4].

However, transferring LLMs to other languages presents several challenges. These models are predominantly developed based on English, posing limitations in transfer learning to other languages [4]. Performance degradation is likely in non-English languages due to linguistic biases and the scarcity of pre-training data [5]. Most LLMs are pre-trained on English data, necessitating additional fine-tuning to achieve specialized performance in other languages [4]. For instance, languages like Korean, which have significant grammatical and lexical differences from English, require additional training to bridge these gaps.

The process of re-training pre-trained models is impractical for many organizations due to the high costs involved. Training large language models from scratch requires substantial computational resources, which are often beyond the reach of small-scale companies and research institutions. Additionally, the data construction process for Supervised Fine-tuning (SFT) and Reinforcement Learning from Human Feedback (RLHF) [6] is time-consuming and expensive. Ensuring the quality of human-annotated data can be particularly challenging, as it requires significant human effort and expertise [7]. This bottleneck is further exacerbated by the need for domain-specific and high-quality datasets for effective language transfer.

Given these challenges, there is a pressing need for efficient and cost-effective methods to transfer LLMs to other languages. Such methods should minimize the computational resources and data requirements while maximizing the performance of the models in the target languages. This paper explores various techniques for achieving this goal, focusing on methods that efficiently enhance performance in languages other than English using minimal data and computational resources.

We compare different approaches, including Supervised Fine-tuning (SFT), Direct Preference Optimization (DPO), and Odds Ratio Preference Optimization (ORPO), and discuss their merits and limitations in the context of cross-lingual transfer learning. By examining these techniques, we aim to provide insights into the most effective strategies for language transfer in LLMs and highlight areas for future research and development.

*Corresponding author

	SFT Wins	Tie	SFT Losses
SFT ($k = 5$) vs. PT ($k = 5$)	41.6	28.5	29.9
SFT ($k = 0$) vs. ORPO ($k = 0$)	21.7	9.6	68.7

Table 1. Comparison of win rates between models optimized with SFT, DPO, ORPO, and baseline models

Model	**BoolQ F1**	**COPA F1**	**HellaSwag F1**	**SentiNeg F1**
PT ($k = 0$)	<u>38.51</u>	<u>57.14</u>	43.29	55.92
SFT ($k = 0$)	37.56	55.38	<u>43.80</u>	<u>57.70</u>
PT ($k = 5$)	45.02	<u>58.30</u>	**52.31**	<u>85.81</u>
SFT ($k = 5$)	<u>62.46</u>	56.53	<u>49.10</u>	80.55

Table 2. Performance on Kobest Benchmark for 0-shot and 5-shot settings

2. Performance Analysis

To evaluate the effectiveness of different language transfer techniques, we conducted experiments on the Kobest natural language processing benchmark. Table 2 presents the 0-shot and 5-shot performance of various models.

2.1 0-shot Performance

For the 0-shot setting, the pre-trained (PT) model achieved an F1 score of 38.51 on BoolQ, 57.14 on COPA, 43.29 on HellaSwag, and 55.92 on SentiNeg. The SFT model, on the other hand, showed slightly lower performance on BoolQ (37.56), COPA (55.38), and HellaSwag (43.80), but slightly better performance on SentiNeg (57.70). This indicates that while SFT can improve performance in some tasks, it may not always outperform the pre-trained model in a 0-shot setting.

2.2 5-shot Performance

In the 5-shot setting, both models showed significant improvements. The pre-trained model achieved an F1 score of 45.02 on BoolQ, 58.30 on COPA, 52.31 on HellaSwag, and 85.81 on SentiNeg. The SFT model outperformed the pre-trained model on BoolQ (62.46) and HellaSwag (49.10) but showed slightly lower performance on COPA (56.53) and SentiNeg (80.55). This suggests that SFT can be particularly effective in scenarios where a small amount of task-specific data is available.

In addition to the Kobest benchmark, we also compared the win rates of different models in various settings. Table 1 presents the win rates for models optimized using SFT, DPO, and ORPO against a baseline model.

3. Conclusion

In this paper, we have examined various techniques for transferring pre-trained Llama-2 models to cross-lingual contexts, focusing on methods such as Supervised Fine-tuning (SFT), Direct Preference Optimization (DPO), and Odds Ratio Preference Optimization (ORPO). Our comparative analysis revealed that each method has its unique strengths and weaknesses, particularly in terms of computational efficiency and performance in different linguistic tasks.

The results from the Kobest benchmark indicate that while SFT generally performs well with additional task-specific data, pre-trained models can still hold their own in zero-shot settings. ORPO, with its resource-efficient approach, offers a promising alternative for organizations with limited computational resources.

4. Limitations

Despite the insights gained, this study has several limitations. Firstly, the impact of varying the ratio of target language to source language data was not thoroughly explored, which could affect performance outcomes. Secondly, the evaluation focused primarily on Korean, leaving the performance in other languages and the source language (English) under-examined. Thirdly, while we used pre-trained models for our experiments, the potential benefits of starting with SFT models were not fully investigated. Lastly, the quality and diversity of the datasets used could influence the generalizability of the findings.

Future research should aim to address these limitations by conducting more comprehensive evaluations across multiple languages and varying data ratios. Additionally, exploring the integration of different pre-trained models and dataset qualities will provide a more holistic understanding of optimizing LLMs for multilingual applications.

5. Acknowledgement

This work was supported by Institute for Information & communications Technology Promotion(IITP) grant funded by the Korea government(MSIT) (RS-2024-00398115, Research on the reliability and coherence of outcomes produced by Generative AI). This research was supported by Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Education(NRF-2021R1A6A1A03045425). This research was supported by Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Education(NRF-2022R1A2C1007616)

Reference

[1] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar *et al.*, "Llama: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.

[2] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale *et al.*, "Llama 2: Open foundation and fine-tuned chat models," *arXiv preprint arXiv:2307.09288*, 2023.

[3] A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Yang, A. Fan *et al.*, "The llama 3 herd of models," *arXiv preprint arXiv:2407.21783*, 2024.

[4] J. Zhao, Z. Zhang, Q. Zhang, T. Gui, and X. Huang, "Llama beyond english: An empirical study on language capability transfer," *arXiv preprint arXiv:2401.01055*, 2024.

[5] M. Muraoka, B. Bhattacharjee, M. Merler, G. Blackwood, Y. Li, and Y. Zhao, "Cross-lingual transfer of large language model by visually-derived supervision toward low-resource languages," *Proceedings of the 31st ACM International Conference on Multimedia*, pp. 3637–3646, 2023.

[6] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray *et al.*, "Training language models to follow instructions with human feedback," *Advances in neural information processing systems*, Vol. 35, pp. 27 730–27 744, 2022.

[7] W. Yuan, R. Y. Pang, K. Cho, S. Sukhbaatar, J. Xu, and J. Weston, "Self-rewarding language models," *arXiv preprint arXiv:2401.10020*, 2024.