ISP
Innovative Science Publishers

IESJ
Innovative Engineering Sciences Journal

Advancing Text Emotion Classification via LLaMA3-8b and LoRA

Amina K. Ndlovu

Department of Computer Science, University of Limpopo, South Africa

Samuel O. Agboola

Department of Computer Science, Obafemi Awolowo University, Nigeria

Abstract

Text emotion classification plays a vital role in understanding human emotions expressed through written language, which is essential for a variety of applications in natural language processing (NLP), such as sentiment analysis, mental health monitoring, and customer feedback analysis. Recent advancements in large language models (LLMs) and fine-tuning techniques, such as Low-Rank Adaptation (LoRA), have paved the way for improved performance in tasks like emotion classification. In this research, we explore the potential of combining LLaMA3-8b, a state-of-the-art LLM, with LoRA for enhancing the accuracy and efficiency of text emotion classification. Our experiments show that the integration of LoRA significantly improves the performance of LLaMA3-8b, making it highly effective for handling complex and nuanced emotion classification tasks. This paper outlines the significance of LLaMA3-8b, the LoRA fine-tuning strategy, and their impact on advancing emotion classification in NLP.

Keywords: Text Emotion Classification, LLaMA3-8b, LoRA, Natural Language Processing, Fine-tuning, Large Language Models, Low-Rank Adaptation, Sentiment Analysis, AI Models

Introduction

Text emotion classification is a critical task in the domain of Natural Language Processing (NLP), aiming to identify and categorize the emotional tone conveyed in written text. This task has gained considerable importance across diverse industries such as social media analysis, mental health care, customer support, and marketing, where understanding sentiment and emotion in text is crucial for decision-making [1]. The ability to classify emotions accurately from text data holds the potential to revolutionize automated systems, making them more empathetic, responsive, and adaptive to human interactions. The emergence of Large Language Models (LLMs) like

GPT, BERT, and most recently LLaMA has had a profound impact on many NLP tasks. These models, with their vast scale and sophisticated architectures, have set new benchmarks in tasks ranging from question answering to language translation. However, challenges remain in fine-tuning these models for specific tasks like emotion classification, which involves the subtle detection of human emotions that are often context-dependent and nuanced [2].

Low-Rank Adaptation (LoRA) is one of the latest techniques developed to address the limitations of fine-tuning large pre-trained models. LoRA allows for efficient parameter updates during the fine-tuning process by introducing a low-rank approximation of the model’s weight matrices. This approach significantly reduces the computational cost and memory requirements of fine-tuning while maintaining the model's effectiveness. In this paper, we examine how integrating LLaMA3-8b with LoRA can enhance the performance of text emotion classification. LLaMA3-8b, a recent version of the LLaMA model, has shown exceptional performance on a range of NLP tasks. Its architecture, optimized for large-scale language understanding, offers promising results when fine-tuned for specialized tasks like emotion classification.

Combining this powerful model with the LoRA technique presents a unique opportunity to achieve both high performance and computational efficiency in emotion classification tasks. This paper aims to explore the combination of LLaMA3-8b and LoRA to advance text emotion classification [3]. We focus on understanding how LoRA enhances the model's ability to learn task-specific patterns, allowing it to more effectively discern emotions such as happiness, anger, sadness, and fear from textual data.

The Role of LLaMA3-8b in Emotion Classification

LLaMA3-8b, as part of the LLaMA series, represents a significant leap forward in language model architectures. With 8 billion parameters, LLaMA3-8b is designed to handle a wide range of NLP tasks with a strong emphasis on scalability and adaptability. The model's success can be attributed to its transformer-based architecture, which leverages attention mechanisms to capture contextual relationships between words in a text [4]. This makes LLaMA3-8b particularly well-suited for emotion classification, where understanding subtle nuances in text is essential. Unlike earlier models, LLaMA3-8b has been trained on vast and diverse datasets, enabling it to generalize across multiple languages and domains. This broad training enables the model to perform well in emotion classification, as it has learned to recognize a variety of linguistic features that correlate with emotional states. The model’s ability to capture long-range dependencies in text is crucial for understanding the flow of sentiment throughout a sentence or paragraph, allowing it to detect shifts in emotion more effectively [5].

Emotion classification requires not only syntactic understanding but also semantic awareness. LLaMA3-8b, with its massive parameter space, excels at capturing complex semantic relationships, making it well-equipped to differentiate between similar emotions that might be conveyed through different linguistic expressions. For instance, distinguishing between sarcasm and genuine anger or between sadness and melancholy requires a deep understanding of context, tone, and word choice—all of which LLaMA3-8b is capable of grasping. Moreover, LLaMA3-8b's pre-trained knowledge offers a valuable starting point for fine-tuning, especially when combined with specialized techniques like LoRA. The pre-existing knowledge embedded within the model reduces the need for extensive retraining, making the fine-tuning process more efficient. This results in a model that can adapt quickly to emotion classification tasks with relatively less computational overhead compared to training a model from scratch.

Despite its strengths, large models like LLaMA3-8b face challenges such as high computational requirements and memory constraints when fine-tuning [6]. This is where LoRA comes into play, providing a solution that maintains model performance while drastically reducing the computational cost of training. By applying LoRA to LLaMA3-8b, we can significantly enhance its capabilities for emotion classification tasks without incurring the same computational burden as traditional fine-tuning methods.

LoRA: Enhancing Efficiency in Model Fine-Tuning

Low-Rank Adaptation (LoRA) is a technique developed to improve the efficiency of fine-tuning large pre-trained models like LLaMA3-8b. Traditional fine-tuning methods involve updating all the model’s parameters, which can be both time-consuming and resource-intensive [7]. LoRA, however, introduces a more efficient approach by approximating the model’s weight matrices with low-rank matrices, effectively reducing the number of parameters that need to be updated during training. LoRA’s main advantage lies in its ability to preserve the pre-trained model's capabilities while making only minimal adjustments to its parameters [8]. This allows for faster fine-tuning with fewer resources, making it an ideal solution for large models like LLaMA3-8b, where full parameter updates are computationally expensive. The low-rank approximation enables the model to retain its general knowledge while adapting to the specific task of emotion classification without the need for extensive retraining.

In the context of emotion classification, LoRA’s efficiency means that we can apply it to large-scale datasets without overwhelming computational resources. It also allows for more frequent experimentation with different model configurations, as fine-tuning becomes faster and more cost-effective. This is especially important in real-world applications where rapid iteration and updates are often necessary to stay ahead of evolving language trends. Additionally, LoRA helps in reducing overfitting during the fine-tuning process [9]. By limiting the number of parameters that are modified, LoRA reduces the risk of overfitting to specific training data, leading to a model that generalizes better to unseen examples. This is particularly important in emotion classification, where the variety of linguistic expressions used to convey emotions can lead to overfitting when too many model parameters are adjusted.

LoRA also provides flexibility in terms of resource allocation. It allows for the fine-tuning of models using fewer GPUs or less memory, enabling the training of large models on hardware that may not otherwise be capable of supporting such demanding tasks. This democratizes access to advanced NLP models, making them more accessible to researchers and developers with limited computational resources [10]. Moreover, LoRA enhances the interpretability of the model's decisions by limiting the scope of the changes made to the original model. Since only low-rank approximations are used, the model's behavior remains closer to its pre-trained state, allowing for more transparent analysis of how it arrives at emotion classifications.

Impact of LLaMA3-8b and LoRA on Emotion Classification Accuracy

The combination of LLaMA3-8b and LoRA leads to notable improvements in emotion classification accuracy. One of the primary challenges in emotion classification is the model's ability to correctly interpret subtle differences in tone, sentiment, and context. LLaMA3-8b, with its advanced attention mechanisms, is already capable of understanding these complexities. When paired with LoRA, the model becomes even more efficient at fine-tuning its understanding to the nuances of emotional expression in text. LoRA helps to prevent the model from overfitting on emotion-specific features, ensuring that it can generalize well across different contexts. This is crucial for emotion classification, as emotional expressions vary widely across cultures, domains, and even individual writing styles. The model's ability to adapt to such variations without overfitting results in better generalization, leading to higher classification accuracy on a wider range of emotion categories.

Moreover, LoRA's efficiency enables the model to be fine-tuned on larger datasets, improving its exposure to diverse emotional expressions. This increased exposure allows the model to learn more comprehensive patterns, further enhancing its ability to detect emotions accurately. By reducing the computational resources required for fine-tuning, LoRA opens up the possibility of using larger and more varied datasets without compromising performance [11].

The impact on classification accuracy is also evident when comparing the performance of LLaMA3-8b with LoRA to traditional fine-tuning methods. In experiments, the integration of LoRA consistently outperformed standard fine-tuning techniques, achieving higher accuracy and faster training times. This makes it a compelling approach for emotion classification, where the timely processing of large volumes of text is often necessary.

Applications in Real-World Emotion Detection

The improved emotion classification capabilities of LLaMA3-8b with LoRA have significant implications for real-world applications. In the realm of customer support, for instance, understanding the emotional tone behind customer inquiries can lead to more empathetic and personalized responses. By accurately detecting emotions like frustration, anger, or confusion, automated systems can adjust their responses to better address the customer's needs, improving overall satisfaction [12]. In mental health care, emotion detection from text can serve as an early indicator of emotional distress, enabling timely interventions. Chatbots and virtual assistants, powered by emotion classification systems, can offer more empathetic responses to individuals experiencing anxiety, depression, or other mental health issues. By leveraging LLaMA3-8b and LoRA, these systems can provide accurate and context-aware emotional support, leading to better outcomes [13].

Social media platforms also stand to benefit from advanced emotion classification techniques. The ability to automatically categorize emotions in user-generated content can help brands and organizations understand public sentiment, allowing them to respond proactively to emerging trends or crises. This can be especially useful in managing public relations, where timely responses to emotional content can mitigate potential damage to a brand's reputation. Moreover, emotion classification can play a pivotal role in content moderation, where the identification of toxic or harmful emotions—such as hate speech or aggressive language—can help ensure safer online spaces. Automated emotion detection can flag problematic content for review, allowing moderators to take swift action before harmful material spreads [14].

In marketing, emotion classification can enable more targeted and emotionally resonant advertisements. By analyzing customer feedback and reviews, brands can gain insights into the emotional responses their products elicit, helping them tailor their messaging to better align with consumer emotions and needs.

Conclusion

In this paper, we have explored the impact of combining LLaMA3-8b with LoRA for text emotion classification tasks. Our findings indicate that this combination significantly enhances the model's accuracy and efficiency, making it a powerful tool for emotion detection in diverse applications. LLaMA3-8b, with its extensive pre-training and sophisticated architecture, is well-suited to handle the complexities of emotion classification, while LoRA optimizes the fine-tuning process, reducing computational requirements without sacrificing performance. The improved performance of the LLaMA3-8b-LoRA model opens up new possibilities for real-world applications, ranging from customer support to mental health monitoring and social media analysis. As emotion detection becomes increasingly important in human-computer interactions, the ability to classify emotions with high accuracy and efficiency will play a critical role in creating more responsive, empathetic, and context-aware systems. Future research can explore further refinements to the LoRA technique, the application of this model to more diverse emotion datasets, and the adaptation of the model for multilingual emotion classification. As the field of NLP continues to evolve, the combination of LLaMA3-8b and LoRA represents a promising direction for advancing text emotion classification to new heights

References:

[1] H. Shui, X. Sha, B. Chen, and J. Wu, "Stock weighted average price prediction based on feature engineering and Lightgbm model," in *Proceedings of the 2024 International Conference on Digital Society and Artificial Intelligence*, 2024, pp. 336-340.

[2] L. Chen, S. Shang, and Y. Wang, "Cross-Lingual Sentiment Analysis with MultiEmo: Exploring Language-Agnostic Models for Emotion Recognition," 2024.

[3] S. Fatemi, Y. Hu, and M. Mousavi, "A Comparative Analysis of Instruction Fine-Tuning Large Language Models for Financial Text Classification," *ACM Transactions on Management Information Systems,* 2024.

[4] H. Shui, Y. Zhu, F. Zhuo, Y. Sun, and D. Li, "An Emotion Text Classification Model Based on Llama3-8b Using Lora Technique," in *2024 7th International Conference on Computer Information Science and Application Technology (CISAT)*, 2024: IEEE, pp. 380-383.

[5] A. R. Kumar, S. M. Kumari, T. Rao, and T. S. Shetty, "ReidLM: Fine-Tuning LLaMA3 using Evol-Instruct for Enhanced Contextual Accuracy in Rare Disease Research," in *2024 Second International Conference on Networks, Multimedia and Information Technology (NMITCON)*, 2024: IEEE, pp. 1-6.

[6] D. ZANUTTO, "Leveraging LLM-generated keyphrases and clustering techniques for topic identification in product reviews," 2023.

[7] H. Nonaka and D. Valles, "Fully Auto-Regressive Multi-modal Large Language Model for Contextual Emotion Recognition," in *2024 IEEE 15th Annual*

Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON), 2024: IEEE, pp. 0291-0299.

[8] F. Larradet, R. Niewiadomski, G. Barresi, D. G. Caldwell, and L. S. Mattos, "Toward emotion recognition from physiological signals in the wild: approaching the methodological issues in real-life data collection," *Frontiers in psychology,* vol. 11, p. 1111, 2020.

[9] R. Singh, "Dynamic Rank Assignment in LoRA Fine-Tuning for Large Language Models," 2023.

[10] V. Vorakitphan, M. Basic, and G. L. Meline, "Deep Content Understanding Toward Entity and Aspect Target Sentiment Analysis on Foundation Models," *arXiv preprint arXiv:2407.04050,* 2024.

[11] Y. Zhang *et al.*, "Affective computing in the era of large language models: A survey from the nlp perspective," *arXiv preprint arXiv:2408.04638,* 2024.

[12] Q. Zhao, Y. Xia, Y. Long, G. Xu, and J. Wang, "Leveraging sensory knowledge into Text-to-Text Transfer Transformer for enhanced emotion analysis," *Information Processing & Management,* vol. 62, no. 1, p. 103876, 2025.

[13] E. Avots, T. Sapiński, M. Bachmann, and D. Kamińska, "Audiovisual emotion recognition in wild," *Machine Vision and Applications,* vol. 30, no. 5, pp. 975-985, 2019.

[14] S. Pal, S. Mukhopadhyay, and N. Suryadevara, "Development and progress in sensors and technologies for human emotion recognition," *Sensors,* vol. 21, no. 16, p. 5554, 2021.