Electronics Letters

LETTER OPEN ACCESS

Fast Fine-Tuning Large Language Models for Aspect-Based Sentiment Analysis

Chaelyn Lee | Jaesung Lee

Department of Artificial Intelligence, Chung-Ang University, Seoul, Republic of Korea

Correspondence: Jaesung Lee (curseor@cau.ac.kr)

Received: 14 March 2024 | **Revised:** 14 May 2025 | **Accepted:** 5 August 2025

Funding: This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) [RS-2021-II211341, Artificial Intelligence Graduate School Program (Chung-Ang University)].

ABSTRACT

The method proposed in this study aims to reduce the execution time required for fine-tuning large language models in aspect-based sentiment analysis. To achieve efficient fine-tuning, the large-language model parameter tuning for new data is accelerated through rank decomposition. Experiments on the SemEval datasets demonstrated that our method consistently outperformed strong baselines such as GPT-ABSA and BART-ABSA across multiple metrics including accuracy, F1-score, precision, and recall while also reducing fine-tuning time by approximately 35%. The experimental results demonstrate a notable decrease in execution time with the proposed approach of the fine-tuning process while preserving the accuracy of polarity prediction.

1 | Introduction

To help increase commercial performance, customer feedback on newly released products is analysed on a daily basis [1]. In the digital age, companies are increasingly reliant on analysing customer feedback regularly to gauge the reception of newly launched products and adjust their market offerings accordingly [2]. Sentiment analysis (SA) studies of sentences or documents typically analyse one paragraph (or one sentence) and summarize only one emotional tendency: positive, negative, or neutral [3]. While enterprises expect an accurate system, the SA systems cannot accurately predict if the sentence does not explicitly express a clear sentiment or an opinion [4]. For this purpose, aspect-based sentiment analysis (ABSA), which aims to predict the polarity of specific entities or aspects of a sentence in customer feedback, can be effectively applied. Conventional methods employ large language models (LLMs) because of their superior capability to understand feedback written in natural language. In practice, pre-trained LLMs are often fine-tuned on target domains, such as appliances, restaurants, or movies, to adapt to domain-specific vocabulary and expression patterns. Among existing approaches, BART-ABSA [5] uses pre-trained BART models to generate end-to-end target sequences without task-specific layers or sub-modules; however, unnecessary parameter updates may occur due to its architecture, originally designed for complex denoising tasks. GPT-ABSA [6] reformulates the ABSA task into a generative sequence modelling problem using a GPT-based language model. Nonetheless, it faces limitations in providing accurate analysis of emerging information due to its slow adaptability to newly introduced data. Customer feedback sentences about new products often contain product-specific terms that frequently change, necessitating continuous and repetitive fine-tuning. However, conventional LLM-based ABSA methods largely overlook this issue, resulting in excessive computational costs for repeated adaptation, which may degrade commercial efficiency. In this paper, we propose a fast fine-tuning strategy based on parameter freezing and rank decomposition to mitigate these limitations [7].

TABLE 1 | Datasets used in the experiments.

Dataset	Train	Positive	Negative	Neutral	Conflict
Lap-2014	1914	818	690	369	37
Res-2014	2980	1744	643	520	73
Res-2016	1289	933	307	49	x

2 | Proposed Method

Given a text dataset $D = \{(S_1, A_1, Y_1), \ldots, (S_{|D|}, A_{|D|}, Y_{|D|})\}$, $S = \langle s_1, s_2, \ldots, s_{|S|} \rangle$ and $A = \{a_1, \ldots, a_{|A|}\}$ represent the sequence of text sentences and set of aspects mentioned in each sentence, respectively. In addition, $Y = \langle y_1, y_2, \ldots, y_{|Y|} \rangle$ represents the sentiment polarity sequence for each sentence and its aspects. The ABSA task of classifying sentiment polarity for aspect terms in a sentence can be formulated as

$$f(S, A) \rightarrow Y, \tag{1}$$

where f is a mapping function that receives the input sentence sequence S and aspect label A and assigns a polarity to each sentence and its aspect. In (1), Y denotes the set of output polarity classes predicted by the model. This function is based on a specific model architecture, such as a generative pre-trained transformer (GPT), which processes input data and extracts information to predict the polarity. In this study, the GPT-2 model was adopted [4].

When models utilizing LLMs for ABSA tasks learn new data, typically, all parameters are fine-tuned, which is the main reason for the high computational cost. To efficiently learn new data, we consider a rank decomposition method that conducts a fine-tuning process by adding pairs of existing weight and trainable rank decomposition matrices in parallel. In this method, the pre-trained weights of the existing neural network are fixed, and only low-rank decomposed weights are learned and added to the existing weights. To clarify this structure, we adopt a low-rank adaptation approach where the full weight matrix is decomposed into a frozen pre-trained matrix Φ_0 and a trainable low-rank residual $\Delta\Phi(\Theta)$. This formulation enables efficient adaptation without updating all parameters. The objective function of the proposed model can be represented as

$$\max_{\Theta} \sum_{(S,A,Y)\in D} \sum_{t=1}^{|Y|} \log(f_{\Phi_0+\Delta\Phi(\Theta)}(y_t | S, a_t, Y_{<t})), \tag{2}$$

where Θ represents the set of parameters for the model to be optimized. The model adjusts these parameters to maximize the log-likelihood of the sentiment polarity sequence for each data point in a given dataset D; $f_{\Phi_0+\Delta\Phi(\Theta)}(y_t | S, a_t, Y_{<t})$ for a given sentence S and aspect a_t. Previous polarity $Y_{<t}$ represents the conditional probability that polarity y_t will occur. Φ_0 is the pre-trained weight matrix and $\Delta\Phi(\Theta)$ denotes the additional weights that are fine-tuned. The sum of these two weights is the final weight used to train the proposed model on new data (Table 1).

TABLE 2 | Experimental results of model performance.

Dataset	Measure	Proposed	GPT-ABSA	BART-ABSA
14lap	Accuracy	**79.37**	78.65	70.09
	F1-score	**81.53**	49.32	70.66
	Precision	**71.00**	70.74	48.51
	Recall	**63.29**	63.15	51.60
14res	Accuracy	**84.84**	84.31	79.23
	F1-score	**85.83**	55.8	77.17
	Precision	68.95	**72.79**	60.88
	Recall	**64.01**	63.64	62.68
16res	Accuracy	**91.56**	89.87	84.79
	F1-score	**92.52**	56.4	73.80
	Precision	78.48	**81.08**	52.15
	Recall	**74.65**	67.51	52.91

3 | Experimental Results

Experiments were performed to prove the effectiveness of the proposed method by comparing two well-known models, BART-ABSA and GPT-ABSA, based on Semeval14 [8] and Semeval16 [9] text datasets, each consisting of laptop and restaurant datasets labelled using aspect terms and sentiment polarity as shown in Table 1: SemEval-2014 Laptop (Lap-2014), SemEval-2014 Restaurant (Res-2014), and SemEval-2016 Restaurant (Res-2016). To train the BART-ABSA model, the learning rate, batch size, and the number of epochs were set to 5×10^{-5}, 5, and 50, respectively; for the GPT2-ABSA model, they were set to 5×10^{-5}, 16, and 50; for GPT2-large, they were set to 3×10^{-2}, 16, and 20, considering the computational cost; and for the proposed method, they were set to 3×10^{-2}, 32, and 20, respectively. Table 2 presents the performance of the proposed method in comparison with BART-ABSA and GPT-ABSA across four evaluation metrics: accuracy, F1-score, precision, and recall. The proposed model outperformed GPT-ABSA and BART-ABSA across all datasets. On 14lap, it achieved the highest accuracy (79.37%) and F1-score (81.53%), demonstrating robust sentiment understanding. On 14res, while GPT-ABSA slightly exceeded in F1 and precision, our model achieved superior recall and accuracy. On 16res, our method showed the most pronounced gains with 91.56% accuracy and 92.52% F1-score, highlighting its effectiveness in handling complex and recent sentiment data. Figure 1 shows the training times of GPT-ABSA and the proposed model using the same LLMs adopted by conventional models: the Semeval14 and Semeval16 datasets. The proposed method also significantly reduced training time compared to GPT-ABSA, achieving up to a 73% reduction on the Restaurant16 dataset (from 25 to 9 min).

4 | Conclusion

In this study, an efficient fine-tuning strategy based on the low-rank decomposition for the LLM-based ABSA method was considered. Experimental results showed that the proposed strategy can yield superior prediction accuracy compared to existing methods. In addition, owing to the fast fine-tuning, the refined

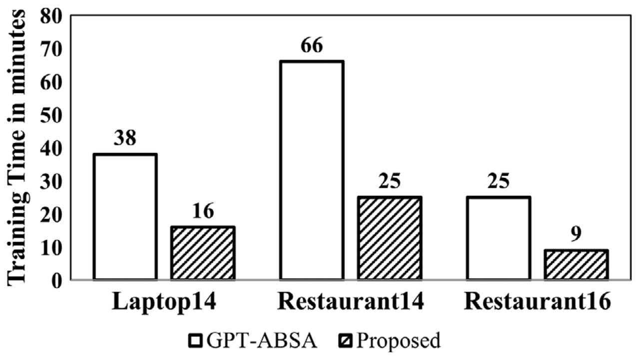

FIGURE 1 | Training time in minutes of the proposed method and conventional model on the Semeval14 and Semeval16 datasets.

model can be obtained quickly, proving its effectiveness in terms of computational efficiency. In future research, we plan to explore changes to the proposed model in search of a more efficient learning method for LLMs.

Author Contributions

Chaelyn Lee: conceptualization, formal analysis, investigation, methodology, writing – original draft. **Jaesung Lee:** conceptualization, funding acquisition, investigation, methodology, project administration, resources, supervision, validation, writing – original draft, writing – review and editing.

Acknowledgements

This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) [RS-2021-II211341, Artificial Intelligence Graduate School Program (Chung-Ang University)].

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The data that support the findings of this study are openly available in SemEval-2014 Task 4, SemEval-2016 Task 5 at https://alt.qcri.org/semeval2014/task4/ , https://alt.qcri.org/semeval2016/task5/index.php?id=data-and-tools.

References

1. V. V. Kumar, K. K. Raghunath, and V. Muthukumaran, "Aspect-Based Sentiment Analysis and Smart Classification in Uncertain Feedback Pool," *International Journal of System Assurance Engineering and Management* 13, no. Suppl 1 (2022): 252–262.

2. M. Wankhade, A. C. S. Rao, and C. Kulkarni, "A Survey on Sentiment Analysis Methods, Applications, and Challenges," *Artificial Intelligence Review* 55, no. 7 (2022): 5731–5780.

3. J. Wang, B. Xu, and Y. Zu, "Deep Learning For Aspect-Based Sentiment Analysis," in *2021 International Conference on Machine Learning and Intelligent Systems Engineering (MLISE)* (IEEE, 2021), 267–271.

4. L. Luc Phan, P. Huynh Pham, and K. Thi-Thanh Nguyen, "SA2SL: From Aspect-Based Sentiment Analysis To Social Listening System For Business Intelligence," in *Knowledge Science, Engineering and Management: 14th International Conference, KSEM 2021* (Springer, 2021), 647–658.

5. H. Yan, J. Dai, and T. Ji, "A Unified Generative Framework For Aspect-Based Sentiment Analysis," in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)* (Association for Computational Linguistics, 2021), 2416–2429.

6. E. Hosseini-Asl, W. Liu, and C. Xiong, "A Generative Language Model for Few-shot Aspect-Based Sentiment Analysis," in *Findings of the Association for Computational Linguistics: NAACL 2022* (Association for Computational Linguistics, 2022), 770–787.

7. E. J. Hu, Y. Shen, and P. Wallis, "LoRA: Low-Rank adaptation of Large Language Models,"*arXiv:2106.09685* (2021).

8. M. Pontiki, D. Galanis, J. Pavlopoulos, H. Papageorgiou, I. Androutsopoulos, and S. Manandhar, "SemEval-2014 Task 4: Aspect Based Sentiment Analysis," in *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval-2014)* (Association for Computational Linguistics, 2014), 27–35.

9. M. Pontiki, D. Galanis, and H. Papageorgiou, "SemEval-2016 Task 5: Aspect Based Sentiment Analysis," in *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)* (Association for Computational Linguistics, 2016), 19–30.