

Optimization techniques for sentiment analysis based on LLM (GPT-3)

Tong Zhan^{1a,5,*}, Chenxi Shi^{1b}, Yadong Shi², Huixiang Li³, Yiyu Lin⁴

^{1a}Computer Science, Columbia University, NY, USA

^{1b}Telecommunication Systems Management, Northeastern University, Boston, MA, USA

²Computer Science, Fudan University, Shanghai, China

³Information Studies, Trine University, AZ, USA

⁴Computer Science and Engineering, Santa Clara University, CA, USA

⁵lubyliu45@gmail.com

*corresponding author

Abstract. With the rapid development of natural language processing (NLP) technology, large-scale pre-trained language models such as GPT-3 have become a popular research object in NLP field. This paper aims to explore sentiment analysis optimization techniques based on large pre-trained language models such as GPT-3 to improve model performance and effect and further promote the development of natural language processing (NLP). By introducing the importance of sentiment analysis and the limitations of traditional methods, GPT-3 and Fine-tuning techniques are introduced in this paper, and their applications in sentiment analysis are explained in detail. The experimental results show that the Fine-tuning technique can optimize GPT-3 model and obtain good performance in sentiment analysis task. This study provides an important reference for future sentiment analysis using large-scale language models.

Keywords: Natural language processing (NLP), Emotion analysis, Large pre-trained language model, GPT-3, Fine-tuning technique.

1. Introduction

Natural language processing (NLP), an important branch of the field of artificial intelligence, aims to enable computers to understand and generate natural language used by humans. With the rapid development of technology, [1] NLP has penetrated into every area of our lives, from intelligent voice assistants to online translation tools to sentiment analysis on social media, its application is everywhere. Sentiment analysis, as one of the important applications of NLP, aims to understand human emotions and attitudes by analyzing emotional colors in text. Humans express emotions in a variety of ways, including text, speech, images, etc. Sentiment analysis focuses on emotional judgment and classification of text data. Sentiment analysis is widely used in social media monitoring, public opinion analysis, product reviews and other fields, which can help enterprises understand users' emotional tendencies, so as to adjust marketing strategies or improve product quality.

However, [2] traditional sentiment analysis methods often rely on hand-designed features and rules, which are difficult to capture complex emotional information in text, and have limited adaptability to

different contexts and cultural backgrounds. In recent years, with the development of deep learning and other technologies, sentiment analysis methods based on large pre-trained language models (such as GPT-3) [3] have gradually become a research hotspot. This paper aims to explore GPT-3-based sentiment analysis optimization technology to improve model performance and effect, and further promote the development of NLP field.

2. Formatting the title, authors and affiliations

2.1. Traditional Natural Language Processing (NLP)

Natural language processing (NLP) is a field of AI that aims to enable computers to understand and use human language in order to perform useful tasks. Natural language processing is divided into two parts: Natural language understanding [4] (NLU) and Natural language generation (NLG). NLP is an area of AI responsible for understanding and processing human language. NLP is part of the overlap between AI, computer science, and linguistics, where the main goal is to make computers understand statements or words expressed in human language.

NLP can be used to solve many problems. Because the amount of data available for text data is very large, it is impossible for people to process all of it. Wikipedia averages 547 new articles per day, and more than 5 million articles in total. Obviously, one person can't read that much information. [5] NLP faces three challenges: collecting data, classifying data, and extracting relevant information.

In addition, voice input methods such as Sogou input method and Baidu input method also play a vital role in them, they can convert speech into text, greatly improving the user's input efficiency. In addition, natural language processing is also widely used in cross-border e-commerce, tourism translation, social media and other fields. However, natural language processing faces a number of technical challenges, the most important of which is the diversity of speech signals. The speech characteristics of different people are different, and the environmental noise will also cause interference to the recognition of speech signals. Secondly, there is the problem of accent and dialect, which vary greatly in different regions and different cultural backgrounds, which brings troubles to natural language processing. In addition, long speech recognition is also a challenge, because the processing of long speech requires higher computing resources and more complex algorithms.

2.2. Sentiment analysis

Sentiment analysis is an important application in the field of artificial intelligence, which can help us understand the emotional color of a piece of text. For example, we can use sentiment analysis to assess the positive or negative tendency of a news report, or understand the public's attitude and emotion towards a certain topic through comments on social media, or extract opinions, analyze themes and dig emotions from text information such as product reviews and movie reviews.

However, due to the variety of emotional expression forms, the design of rules needs to take into account different emotional expression ways, and this method is not as robust and adaptable as other methods. An emotion dictionary is a dictionary of emotion words in which each word is labeled positive, negative, or neutral, so that the emotion tendency of the text can be obtained by counting the number and polarity of emotion words in the text. Traditional machine learning methods include naive Bayes, support vector machines [6-7](SVM) and random forest algorithms. The basic idea of these methods is to use the existing annotation data to train the classifier, and then use the trained classifier to perform sentiment analysis on the new text. In this way, although a large number of labeled data can be used for training, thereby improving the performance of the classifier. However, features need to be designed manually, and for different tasks and data sets, features need to be redesigned, and the training process is time-consuming.

Deep learning methods include models such as recurrent neural networks (RNN), short term memory networks (LSTM), and convolutional neural networks [8] (CNN). Among them, RNNs and LSTMs can deal with sequence data well, while CNNs can deal with local information in text effectively. These models typically include an embedding layer to translate the text into a vector representation, one or

more hidden layers to extract features, and finally an output layer to predict the emotional polarity of the text. However, most of the time, a large amount of labeled data is needed to train the model, and the training process is complicated and time-consuming.

2.3. Large language model GPT-3

Large language model (LLMs) is a special class of pre-trained language model (PLMs), which is obtained by expanding the model size, pre-trained corpus and computational power. LLMs exhibit special capabilities due to their enormous size and pre-training on large amounts of text data, allowing them to achieve excellent performance in many natural language processing tasks without any task-specific training. The era of [6] LLMs began with [9] OpenAI's GPT-3 model and grew exponentially in popularity after the introduction of models such as ChatGPT and GPT4. We refer to GPT-3 and its successor OpenAI models (including ChatGPT and GPT4) [10] as the GPT-3 Family of Large Language Models (GLLMs).

With the growing popularity of GLLMs, especially in the research community, there is an urgent need for a comprehensive review that summarizes recent research advances across multiple dimensions to provide insights into future research directions. In our review paper, we first introduce basic concepts such as transformers, transfer learning, self-supervised learning, pre-trained language models, and large language models. In conclusion, this comprehensive review paper will be a good resource for those in academia and industry to learn about the latest research related to the GPT-3 family of large language models. Indexed terms - Large Language Model, GPT-3, ChatGPT, GPT-4, transformer.

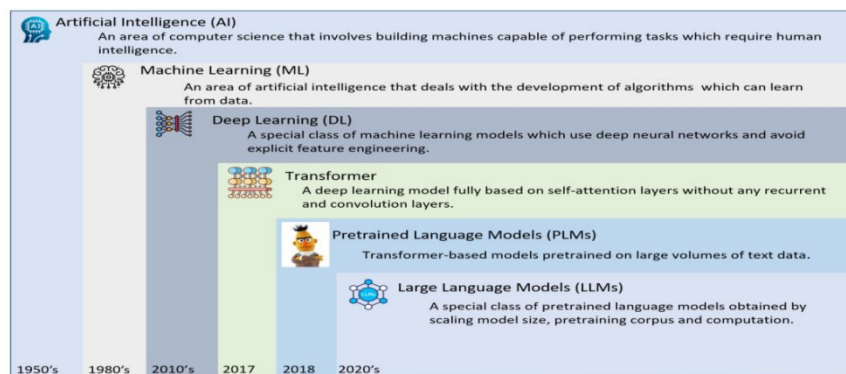


Figure 1. Evolution history of large language models

LLMs leverages contextual learning (ICL), a new learning paradigm that does not require task-specific fine-tuning and a large number of labeled instances [11]. LLMs treats any NLP task as a conditional text generation problem and generates the required text output based solely on input prompts, including a task description, test input, and optionally some examples. Figure 1 shows the evolution of AI from machine learning to large language models.

3. Methodology

3.1. GPT-3 model

GPT-3 Language Models are Few Shot Learners Grand model is a very large scale language model developed by OpenAI [12]. The GPT-3 Grand model has a staggering parameter scale of 175 billion parameters, making it one of the largest models of its kind. The GPT-3 big model is based on the Transformer architecture, which is a very popular architecture in natural language processing. It uses a self-attention mechanism, which enables the model to automatically associate information from different locations and take context into account when generating text. This mechanism helps the model better understand and generate text that conforms to syntactic and semantic rules.

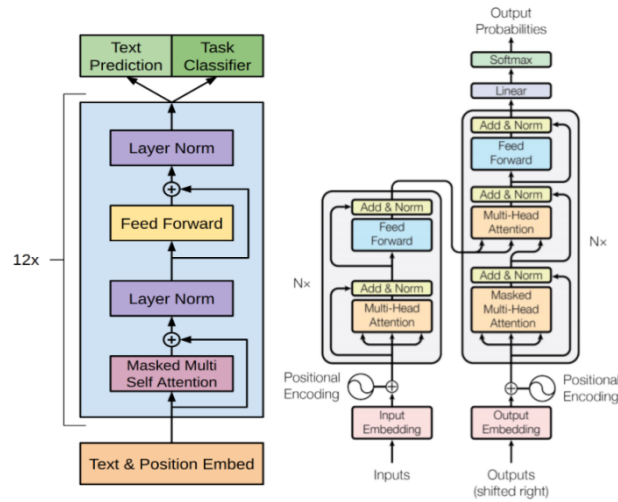


Figure 2. The Transformer - model architecture.

In Figure 2, we describe in detail the basis of GPT-3's network architecture. GPT-3 uses a model architecture called Transformer, which has been hugely successful in the field of natural language processing. The Transformer model consists of multiple stacked self-attention layers, each consisting of two sub-layers: multi-head attention and feedforward neural networks. The self-attention mechanism enables the model to maintain long-distance dependencies when processing sequence data, which is particularly important for natural language processing tasks. In addition, GPT-3 introduces a large number of parameters and a deeper network structure, allowing the model to capture more complex linguistic patterns and semantic information.

3.2. Fine-tuning technique

The process of fine-tuning is not a simple splicing of a [13] pre-trained model with a domain-specific dataset, but a process of fine tuning. First, we need to select the appropriate hyperparameters to adjust the model, such as learning rate, batch size, and training rounds. The learning rate controls the rate at which model parameters are updated, the batch size affects the number of samples taken for each update, and the training rounds determine how many times the model is trained on the entire data set. By adjusting these hyperparameters properly, the model can converge faster and achieve better performance during fine tuning.

To address these issues, we can mitigate the underfitting problem by preventing overfitting through appropriate regularization techniques such as Dropout and data enhancement methods, while increasing the complexity of the model to improve its expressiveness. By taking these factors into consideration, the fine-tuning process can be optimized effectively, and the performance and generalization ability of the model on specific tasks can be improved.

3.3. Application of Transfer Learning

Transfer Learning (Transfer Learning) is to use the learned and trained model parameters as the starting parameters of the new training model. Transfer learning is a very important and commonly used strategy in deep learning.

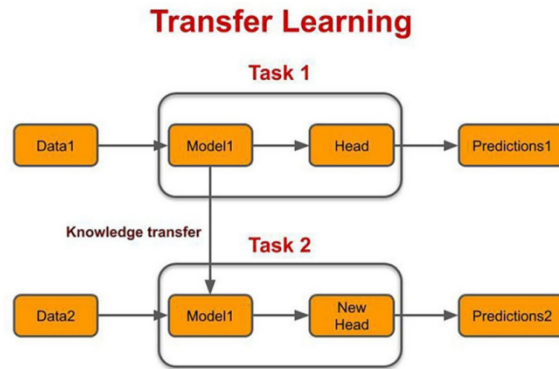


Figure 3. Transfer Learning architecture

At present, Transfer Learning is mainly used in the field of computer and causal inference. It focuses on how models learned from one population or task can be transferred to new target populations or scenarios[14]. In computer science, Transfer Learning is A machine learning method that takes the model developed for task A as the initial point and re-uses it in the process of developing the model for Task B. Transfer learning is a new task that improves learning by transferring knowledge from a related task that has already been learned, and while most machine learning algorithms are designed to solve a single task, the development of algorithms that facilitate transfer learning is an ongoing topic of interest in the machine learning community. Sometimes, the observed covariates in the experiment are often less than the number of covariates in the observed target population, so the application of the above method has obvious limitations. In the computer field, there are two very important concepts in transfer learning: [14-16] Domain and Task. domain can be understood as a specific domain at a certain time, for example, book review and TV series review can be regarded as two different domains, and tasks are things to be done, such as sentiment analysis and entity recognition are two different tasks.

3.4. Sentiment analysis optimization

GPT-3 has a staggering 175 billion parameters scale and has been pre-trained on a large amount of text on the Internet, giving it powerful language understanding and generation capabilities. Among them, Fine-tuning technology is one of the important functions of GPT-3 and other pre-trained models. Fine-tuning improves the performance and adaptability of the model by training it specifically to the needs of a specific task or domain.

Data preparation and preprocessing are important steps in sentiment analysis model training. First, we obtained online data on hotel reviews, including the content of the reviews and the corresponding emotional labels (1 for positive, 0 for negative). We then need to convert this data into a format acceptable to the GPT model, i.e. {"prompt": "<prompt text>", "completion": "<ideal generated text>"}. To validate and format the data, we use tools provided by openai and operate from the command line. Along the way, we added a suffix separator and a space character at the beginning of the finish to help the model better understand the input data. At the same time, we chose to split the data into training sets and validation sets in order to train and evaluate the model. In the end, the process produces two jsonl files that contain the formatted data and provide reference commands and an estimated time for training.

3.5. Model result

Before training the model, we first add the API key to the environment variable to make calls to the openai API. We then use the command line tools provided by openai to create the fine-tuning task. When creating the fine-tuning task, we specified the file paths for the training set and validation set, and used the --compute_classification_metrics parameter to calculate the model's performance metrics for the classification task, including accuracy, precision, recall, and F1 score. At the same time, we specify the positive class or positive example in the classification task through the --classification_positive_class

parameter. However, when executing the command, the model parameters may have been inadvertently not specified, resulting in the Curie model being selected by default for fine tuning. To cancel this training, we used the commands provided by openai to cancel the fine-tuning task. By querying the list of fine-tuning tasks, we found that the model chosen by default is Curie.

Next, we resumed the fine-tuning task and waited for the training to complete. If the training process is interrupted unexpectedly, we can use the commands provided by openai to track the progress of the fine-tuning task. After the training is complete, we can view the performance metrics of the model by querying the results of the fine-tuning task. According to the results, the accuracy of the model is 0.85, which indicates that the model has achieved good performance in the classification task.

3.6. *Experimental discussion*

In the experimental discussion, we examine the performance of the GPT-3 model and the impact of Fine-tuning for sentiment analysis optimization. Leveraging its extensive parameter scale and advanced language processing capabilities, GPT-3, based on the Transformer architecture, incorporates self-attention mechanisms for text understanding and generation. Through Fine-tuning, we adeptly adjust model parameters to better suit emotion analysis tasks, enhancing overall performance. Pre-processing hotel review data and formatting it for GPT-3, we validate, split, and optimize the model for sentiment analysis. With a fine-tuned accuracy of 0.85, our results demonstrate the model's efficacy. We apply the trained model to real sentiment analysis tasks via API calls, assessing both model performance metrics and practical application outcomes. These experiments elucidate Fine-tuning's role in sentiment analysis optimization, facilitating further discussions on model refinement and application potential, thus providing valuable insights for future research and deployment.

4. Conclusion

The future trajectory of GPT-3 and its successors in natural language processing appears promising despite existing challenges. As computing resources evolve and technology advances, we anticipate significant performance enhancements and efficacy improvements. Upgrades in hardware infrastructure will likely bolster the efficiency of GPT-3 during training and reasoning, broadening its applicability across diverse scenarios.

Moreover, advancements in sentiment analysis optimization will be a focal point of future research and practice. Beyond Fine-tuning, researchers will explore alternative techniques like transfer learning and model distillation to refine the model's performance. Integrating domain knowledge and contextual information promises to enhance the model's understanding and analytical capabilities further.

Expanding beyond natural language processing, GPT-3 holds potential for diverse applications, including intelligent decision-making assistance, customer service systems, and creative assistants. Despite ongoing technical hurdles, the continuous evolution of technology and the widening scope of applications suggest that GPT-3 and its successors will play an increasingly vital role, driving progress in natural language processing and contributing to societal advancement.

References

- [1] Chowdhary, KR1442, and K. R. Chowdhary. "Natural language processing." *Fundamentals of artificial intelligence* (2020): 603-649.
- [2] Nadkarni, Prakash M., Lucila Ohno-Machado, and Wendy W. Chapman. "Natural language processing: an introduction." *Journal of the American Medical Informatics Association* 18.5 (2011): 544-551.
- [3] Fanni, S. C., Febi, M., Aghakhanyan, G., & Neri, E. (2023). Natural language processing. In *Introduction to Artificial Intelligence* (pp. 87-99). Cham: Springer International Publishing.
- [4] Gillioz, A., Casas, J., Mugellini, E., & Abou Khaled, O. (2020, September). Overview of the Transformer-based Models for NLP Tasks. In *2020 15th Conference on Computer Science and Information Systems (FedCSIS)* (pp. 179-183). IEEE.

- [5] Choudhury, M., Li, G., Li, J., Zhao, K., Dong, M., & Harfoush, K. (2021, September). Power Efficiency in Communication Networks with Power-Proportional Devices. In 2021 IEEE Symposium on Computers and Communications (ISCC) (pp. 1-6). IEEE.
- [6] Ethayarajh, K. (2019). How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings. arXiv preprint arXiv:1909.00512.
- [7] Zad, S., Heidari, M., James Jr, H., & Uzuner, O. (2021, May). Emotion detection of textual data: An interdisciplinary survey. In 2021 IEEE World AI IoT Congress (AIIoT) (pp. 0255-0261). IEEE.
- [8] Zhou, Tong. "Improved sales forecasting using trend and seasonality decomposition with lightgbm." 2023 6th International Conference on Artificial Intelligence and Big Data (ICAIBD). IEEE, 2023.
- [9] Nguyen, Dat Quoc, Thanh Vu, and Anh Tuan Nguyen. "BERTweet: A pre-trained language model for English Tweets." arXiv preprint arXiv:2005.10200 (2020).
- [10] Min, Bonan, et al. "Recent advances in natural language processing via large pre-trained language models: A survey." ACM Computing Surveys 56.2 (2023): 1-40.
- [11] Ma, Haowei, Cheng Xu, and Jing Yang. "Design of Fine Life Cycle Prediction System for Failure of Medical Equipment." Journal of Artificial Intelligence and Technology 3.2 (2023): 39-45.
- [12] Gunel, B., Du, J., Conneau, A., & Stoyanov, V. (2020). Supervised contrastive learning for pre-trained language model fine-tuning. arXiv preprint arXiv:2011.01403.
- [13] Barone, Antonio Valerio Miceli, et al. "Regularization techniques for fine-tuning in neural machine translation." arXiv preprint arXiv:1707.09920 (2017).
- [14] Renda, Alex, Jonathan Frankle, and Michael Carbin. "Comparing rewinding and fine-tuning in neural network pruning." arXiv preprint arXiv:2003.02389 (2020).
- [15] Vrbančič, G., & Podgorelec, V. (2020). Transfer learning with adaptive fine-tuning. IEEE Access, 8, 196197-196211.
- [16] Floridi, Luciano, and Massimo Chiriatti. "GPT-3: Its nature, scope, limits, and consequences." Minds and Machines 30 (2020): 681-694.