



Leveraging Transformer Neural Networks for Enhanced Sentiment Analysis on Online Platform Comments

Jiayi Tang

University of California, Davis, California, USA

jyitang@ucdavis.edu

ABSTRACT

The relentless surge in the volume of data generated across online platforms necessitates sophisticated approaches to glean actionable insights, particularly in discerning the sentiment encased within user-generated comments. This paper presents an intricate exploration of applying deep learning technologies, specifically the Transformer neural network model, in classifying sentiments expressed in a copious corpus of 1.6 million tweets from the Sentiment140 dataset. Endeavoring beyond conventional models such as Support Vector Machines (SVM) and Long Short-Term Memory networks (LSTM), this research meticulously evaluates the performance enhancements offered by the Transformer architecture in the realm of sentiment analysis. Our methodology entails the comprehensive utilization of the Transformer model, renowned for its prowess in handling sequential data, to unearth the underlying sentiments in the tweets, categorizing them aptly as positive, neutral, or negative. The empirical results emanating from this study delineate a noteworthy improvement in classification accuracy and model robustness, solidifying the Transformer's position as a formidable architecture for sentiment analysis tasks. The enlightening findings of this research offer a pivotal contribution to the continuous efforts in augmenting sentiment analysis methodologies, emphasizing the substantial potentials of integrating advanced deep learning models in processing and analyzing online platform comments. Through this work, we envisage fostering innovative advancements in the field, cultivating enhanced techniques for accurate, efficient, and impactful sentiment analysis in the burgeoning digital landscapes.

CCS CONCEPTS

• **Computing methodologies** → Machine learning; Machine learning approaches.

KEYWORDS

Transformer Neural Networks, Sentiment Analysis, Deep Learning, Online Platform Comments, Classification, NLP, LSTM

ACM Reference Format:

Jiayi Tang. 2023. Leveraging Transformer Neural Networks for Enhanced Sentiment Analysis on Online Platform Comments. In *2023 4th International Conference on Machine Learning and Computer Application (ICMLCA 2023)*, October 27–29, 2023, Hangzhou, China. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3650215.3650260>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ICMLCA 2023, October 27–29, 2023, Hangzhou, China

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0944-9/23/10

<https://doi.org/10.1145/3650215.3650260>

1 INTRODUCTION

In the evolving realms of digital communication, online platforms have burgeoned as powerful mediums where users actively share opinions, experiences, and feedback through comments or reviews. Such platforms brim with rich, user-generated content that encapsulates public sentiments on a plethora of topics ranging from products and services to global events. The immense volume and velocity at which this data is generated necessitate the advent of sophisticated sentiment analysis methodologies to accurately decipher and categorize the underlying emotions and opinions [1–3].

Sentiment analysis, or opinion mining, stands as a pivotal facet of natural language processing (NLP) and text analytics, where the goal orbits around determining the sentiments harbored in textual content [4, 5]. This domain, fueled by the exponential growth of online interactions, faces the continuous challenge of adapting to the nuanced and evolving language used in digital communication. This dynamism has propelled the advancement of techniques that not only discern the polarity of sentiments but also comprehend the subtleties and complexities inherent in human language.

In this research, the spotlight is cast on enhancing sentiment analysis through the innovative application of deep learning technologies for classifying comments retrieved from online platforms. A meticulous exploration is conducted utilizing the Sentiment140 dataset, a robust collection featuring 1.6 million tweets, thereby presenting a vast and diverse playground for model evaluation and analysis [6, 7]. In pursuit of excellence in classification performance, this study embarks on employing the Transformer neural network model, steering away from traditional architectures such as Support Vector Machines (SVM) and Long Short-Term Memory networks (LSTM) [8–10]. The Transformer model, revered for its capacity to process sequential data with remarkable efficiency and accuracy, unveils new horizons in sentiment analysis capabilities. Its unique attention mechanism enables the model to focus on different parts of the input sequence, enhancing its ability to grasp contextual nuances and implicit meanings, which are crucial in accurately interpreting sentiments.

This paper unfolds with a thorough evaluation of the proposed Transformer-based model, juxtaposing its performance against prevailing models like SVM and LSTM. Through systematic experimentation and rigorous analysis, the research reveals enlightening insights into the superiority of the Transformer model in enhancing classification accuracy and robustness in sentiment analysis tasks. The findings not only underscore the technological advancements in NLP but also highlight the potential of these innovations in reshaping the landscape of sentiment analysis. Thus, this exploration stands as a significant stride towards unveiling novel methodologies and fostering innovation in sentiment analysis practices within the dynamic online landscapes.

2 RELATED WORK

Sentiment analysis, a crucial branch of natural language processing (NLP), has garnered substantial attention in the realm of computational linguistics and text mining. Its essence lies in discerning subjective information, emotions, or opinions from text, providing valuable insights across diverse domains such as marketing, politics, and social sciences. Several methodologies and algorithms have been proposed in literature, aiming at the efficient classification of sentiments expressed in textual data. Traditional methods employed in sentiment analysis predominantly include machine learning algorithms such as Support Vector Machines (SVM), Naive Bayes, and Random Forests [11–13]. SVM, in particular, has been a popular choice owing to its effectiveness in high-dimensional spaces, and has been widely utilized for binary as well as multi-class sentiment classification tasks [14, 15].

Deep learning has ushered in a renaissance in NLP and sentiment analysis. Among the neural network architectures, Long Short-Term Memory networks (LSTM) have been a noteworthy contribution. LSTMs, with their capability to capture long-term dependencies in sequential data, have proven to be particularly potent for tasks involving sequential data such as text and time-series [16, 17]. However, the evolution in the field continued with the emergence of the Transformer model, which has revolutionized the landscape, originally proving its mettle in machine translation and subsequently across various NLP tasks [18, 19]. It eschews recurrence mechanisms in favor of self-attention mechanisms, facilitating parallel processing of sequences and a global view of input data.

Remarkably, Transformer models have not confined their prowess to NLP. They have also made significant inroads into computer vision [20–22]. Demonstrating remarkable flexibility and effectiveness in handling image data, they have facilitated groundbreaking performance in tasks like image classification and object detection. This versatility indicates a paradigm shift in the approach to processing and analyzing data, irrespective of its nature.

Furthermore, recent advancements in this domain have led to the development of models like BERT (Bidirectional Encoder Representations from Transformers) and GPT (Generative Pre-trained Transformer), which have taken the capabilities of Transformers a notch higher. These models, pre-trained on vast amounts of textual data, have shown exceptional performance in a variety of NLP tasks, including sentiment analysis, by capturing contextual meanings more effectively than ever before [23, 24].

In synthesis, while traditional machine learning and LSTM have paved the way in sentiment classification, the Transformer model stands out as a versatile and powerful architecture, heralding unprecedented advancements not only in NLP but extending its remarkable capabilities into the realms of computer vision as well. This study seeks to build upon this rich tapestry of related work, aiming to further elucidate the capabilities of Transformer models in sentiment analysis and classification tasks, while emphasizing the continuous evolution and integration of AI technologies across various domains of data analysis.

3 METHODOLOGY

In this section, we elucidate our proposed methodology that seamlessly integrates local-global self-attention mechanisms, enhanced multi-layer perceptrons (MLP), and sophisticated encoding techniques, as illustrated in Figure 1. This holistic approach aims to bolster the Transformer’s capacity to effectively discern and classify sentiments within the tweets encompassed in the Sentiment140 dataset.

3.1 Encoding Technique

Our model initiation involves a robust encoding process. We utilize pre-trained word embeddings as a foundation, where each token in the input sequence is mapped to a high-dimensional continuous vector space [23, 24]. Additionally, we employ positional encoding to infuse the model with the requisite awareness of the sequential order of words in a tweet. Mathematically, the encoding can be represented in Equation 1) as:

$$E(x) = W_e \cdot x + P \quad (1)$$

where $E(x)$ denotes the encoded input, W_e represents the word embeddings, and P stands for the positional encoding.

3.2 Local-Global Self-Attention

We enhance the standard self-attention mechanism in the Transformer model by integrating a local-global attention strategy [25, 26]. This novel approach (seen in Equation 2) allows each position in the input sequence to focus on a restricted window of surrounding positions for local context and the entire sequence for global context simultaneously.

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (2)$$

where Q, K , and V are queries, keys, and values, respectively. d_k is the dimension of the keys.

3.3 Improved Multi-Layer Perceptron

In our ongoing quest for model refinement and enhancement, we turn our focus towards the optimization of the Multi-Layer Perceptron (MLP) housed within the Transformer architecture. This fine-tuning aims to bolster the model’s capacity to effectively understand and process the intricacies of sentiment in textual data, fostering a more nuanced and accurate classification. An intricate part of our model, the MLP, is innovatively modified to encompass a succession of fully connected layers, each characterized by activation functions and layer normalization processes. The enhanced design of the MLP is succinctly encapsulated in Equation 3):

$$\text{MLP}(x) = \text{LayerNorm}(x + \text{ReLU}(W_2 \cdot \text{LayerNorm}(W_1 \cdot x + b_1) + b_2)) \quad (3)$$

where W_1 and W_2 denote weight matrices, and b_1 and b_2 are bias terms.

3.4 Model Architecture

Our model architecture fosters a hierarchical construction where multiple layers of the enhanced Transformer blocks are stacked. Each block is enriched with our local-global self-attention and




Figure 1: The Proposed Sentiment Classification Network with Local-Global Self Attention Mechanism and Enhanced MLP




Figure 2: The Distribution of Sentiment Analysis Dataset

improved MLP, facilitating nuanced extraction and processing of features relevant to sentiment classification.

3.5 Loss Function and Optimization

We employ a categorical cross-entropy loss function (seen in Equation 4), conducive to the multi-class classification nature of our task. An Adam optimizer is used to minimize the loss function, fine-tuning the model parameters iteratively to enhance performance [27].

$$L = - \sum_{i=1}^N y_i \log(p_i) \quad (4)$$

where y_i are the true labels, and p_i denotes the model's predicted probabilities. Through the strategic amalgamation of the enhanced encoding technique, local-global self-attention, and improved MLP in our methodology, we aim to cultivate a model that exhibits heightened sensitivity and accuracy in sentiment classification on the substantial Sentiment 140 dataset.

4 EXPERIMENT

4.1 Dataset and Preprocessing

We employed the Sentiment140 dataset, which encompasses 1.6 million tweets, as the bedrock for our experiment. The distribution of the tweets as to whether positive or negative is briefly sketched in Figure 2. The tweets were preprocessed to expunge noise and enhance data quality, which included the removal of special characters, URLs, and user mentions. Each tweet was tokenized, and common stopwords were removed to enrich the dataset with meaningful content for analysis. Additionally, the dataset was augmented with sentiment-specific preprocessing steps, such as emoticon and slang

interpretation, to preserve and interpret sentiment-laden expressions often found in social media text.

The preprocessing also involved transforming hashtags into usable terms, as they often contain key sentiment indicators. This transformation was achieved by employing a camel case splitter, which efficiently converts concatenated words in a hashtag into separate terms without losing their contextual meaning. Furthermore, to address the imbalanced nature of the dataset, we applied techniques like oversampling for minority classes to ensure a fair representation of all sentiment categories in the training process.

We split the dataset as 70% for training, 15% for validation, and the rest 15% for testing. This distribution was chosen to provide a substantial amount of data for training the model while allowing for comprehensive validation and testing, ensuring the reliability and robustness of our sentiment analysis model.

4.2 Model Configuration

Our proposed model, enriched with local-global self-attention and improved MLPs, was configured meticulously. The model comprises multiple layers of enhanced Transformer blocks, each harnessing the power of our refined attention mechanism and MLP structure. Hyperparameters, such as learning rate, batch size, and number of epochs, were tuned through preliminary experiments to identify configurations that yield optimal performance.

The number of Transformer layers, along with the dimensionality of the input and output vectors and the size of the attention heads, were calibrated to balance computational efficiency with model complexity. Additionally, dropout layers were introduced as a regularization technique to prevent overfitting, especially given the model's exposure to a large volume of data.

To enhance the model's capacity to handle the nuances of sentiment analysis, we incorporated an adaptive learning rate scheduler, which adjusts the learning rate based on the model's performance over time. This approach ensures that the model converges to an optimal solution more effectively, reducing the likelihood of getting trapped in local minima. The model's performance was further optimized by employing advanced techniques such as gradient clipping and batch normalization, which not only improved the training stability but also enhanced the overall learning process. Gradient clipping was particularly crucial in managing the infamous problem of exploding gradients, often encountered in deep neural networks.

In conclusion, the rigorous model configuration, combined with sophisticated preprocessing and hyperparameter tuning, establishes a robust foundation for our sentiment analysis model, designed to excel in accuracy and efficiency in classifying sentiments from the vast and varied data presented in the Sentiment140 dataset.




Figure 3: The Confusion Matrix of Our Proposed Model

4.3 Baseline Models

For a comprehensive evaluation, our model was juxtaposed against several baseline models, including traditional machine learning algorithms such as SVM and deep learning architectures like LSTM [28–32]. Additionally, to broaden the spectrum of our comparative analysis, we included other relevant models like Bidirectional LSTM (BiLSTM), recurrent neural networks (RNN) with varying layers, Bernoulli Naive Bayes, and Linear Support Vector Classification (LinearSVC). This diverse array of models provided a rich context for assessing the effectiveness and improvements manifested by our proposed methodology. By comparing against both traditional and advanced techniques, our aim was to demonstrate not only the efficacy of our model in the context of current standards but also its potential to set new benchmarks in sentiment analysis.

4.4 Results and Discussion

Our experimental results revealed that the proposed model exhibited significant improvements over the baseline models across various evaluation metrics. The confusion matrix of our proposed model is briefly illustrated in Figure 3. The integration of local-global attention mechanisms and enhanced MLP structures within the Transformer architecture was instrumental in bolstering the model’s capacity to accurately classify sentiments within the tweets, and the quantitative results are illustrated in Table 1.

5 CONCLUSION

This research heralds a significant stride forward in sentiment analysis, particularly in classifying sentiments in tweets from online platforms. Through meticulous experimentation, our proposed model, infused with enhanced local-global self-attention mechanisms and optimized MLP structures, showcased remarkable prowess in handling the multifaceted Sentiment140 dataset, outperforming several

baseline models including SVM, LSTM, BiLSTM, and various configurations of RNNs.

The experimental outcomes underscore the model’s enhanced capacity to accurately discern and classify sentiments, advocating for the efficacy of our proposed enhancements within the Transformer architecture. This study thus contributes valuable insights and advancements to the field of sentiment analysis, paving the way for further innovations and refinements in employing deep learning technologies for nuanced and effective sentiment classification in online textual data.

Future work may explore further optimizations and adaptations of the model to extend its applicability and performance across various domains and datasets. Additionally, there is scope for exploring the integration of multimodal data, such as combining textual analysis with audio and visual cues, to enhance the depth and accuracy of sentiment analysis. Finally, addressing challenges such as model interpretability and bias mitigation remains a crucial area for ongoing research, ensuring that sentiment analysis tools are not only powerful but also fair and transparent in their applications.

REFERENCES

- [1] Medhat, Walaa, Ahmed Hassan, and Hoda Korashy. "Sentiment analysis algorithms and applications: A survey." *Ain Shams engineering journal* 5.4, 2014, 1093-1113.
- [2] Hussein, Doaa Mohey El-Din Mohamed. "A survey on sentiment analysis challenges." *Journal of King Saud University-Engineering Sciences* 30.4, 2018, 330-338.
- [3] Devika, M. D., C*Sunitha, and Amal Ganesh. "Sentiment analysis: a comparative study on different approaches." *Procedia Computer Science* 87, 2016, 44-49.
- [4] Kolchyna, Olga, *et al.* "Twitter sentiment analysis: Lexicon method, machine learning method and their combination." *arXiv preprint arXiv:1507.00955*, 2015.
- [5] Gupta, Bhumika, *et al.* "Study of Twitter sentiment analysis using machine learning algorithms on Python." *International Journal of Computer Applications* 165.9, 2017, 29-34.
- [6] Behera, Rabi Narayan, Manan Roy, and Sujata Dash. "A Novel Machine Learning Approach for Classification of Emotion and Polarity in Sentiment140 Dataset." *Conference: 3rd International Conference on Business & Information Management*. 2016.
- [7] Heredia, Brian, *et al.* "Cross-domain sentiment analysis: An empirical investigation." *2016 IEEE 17th International Conference on Information Reuse and Integration*. IEEE, 2016.
- [8] Vaswani, Ashish, *et al.* "Attention is all you need." *Advances in neural information processing systems* 30, 2017.
- [9] Moraes, Rodrigo, João Francisco Valiati, and Wilson P. Gavião Neto. "Document-level sentiment classification: An empirical comparison between SVM and ANN." *Expert Systems with Applications* 40.2, 2013, 621-633.
- [10] Rao, Guozheng, *et al.* "LSTM with sentence representations for document-level sentiment classification." *Neurocomputing* 308, 2018, 49-57.
- [11] Wang, Jenq-Haur, *et al.* "An LSTM approach to short text sentiment classification with word embeddings." *Proceedings of the 30th conference on computational linguistics and speech processing*. 2018.
- [12] Liu, Yang, Jian-Wu Bi, and Zhi-Ping Fan. "A method for multi-class sentiment classification based on an improved one-vs-one (OVO) strategy and the support vector machine (SVM) algorithm." *Information Sciences* 394, 2017, 38-52.
- [13] Al Amrani, Yassine, Mohamed Lazaar, and Kamal Eddine El Kadiri. "Random forest and support vector machine based hybrid approach to sentiment analysis." *Procedia Computer Science* 127, 2018, 511-520.
- [14] Pu, Xiaojia, Gangshan Wu, and Chunfeng Yuan. "Exploring overall opinions for document level sentiment classification with structural SVM." *Multimedia Systems* 25, 2019, 21-33.
- [15] Luo, Fang, Cheng Li, and Zehui Cao. "Affective-feature-based sentiment analysis using SVM classifier." *2016 IEEE 20th International Conference on Computer*

Table 1: The performance of our proposed model against baseline methods on test set.

Model	LSTMx2	LSTMx3	SVM	BiLSTM	RNNx3	RNNx5	BernoulliNB	LinearSVC	Ours
Accuracy	0.750	0.765	0.683	0.743	0.721	0.764	0.801	0.785	0.821

- Supported Cooperative Work in Design. IEEE, 2016.
- [16] Jin, Ning, *et al.* "Multi-task learning model based on multi-scale CNN and LSTM for sentiment classification." IEEE Access 8, 2020, 77060-77072.
 - [17] Huang, Minlie, Yujie Cao, and Chao Dong. "Modeling rich contexts for sentiment classification with lstm." arXiv preprint arXiv:1605.01478. 2016.
 - [18] Brauwiers, Gianni, and Flavius Frasinca. "A survey on aspect-based sentiment classification." ACM Computing Surveys 55.4, 2022, 1-37.
 - [19] Yin, Da, Tao Meng, and Kai-Wei Chang. "Sentibert: A transferable transformer-based architecture for compositional sentiment semantics." arXiv preprint arXiv:2005.04114, 2020.
 - [20] Wang, and *et al.* "Adversarial Vision Transformer for Medical Image Semantic Segmentation with Limited Annotations." 2022.
 - [21] Wang, ZY, Nanqing Dong, and Irina Voiculescu. "Computationally-efficient vision transformer for medical image semantic segmentation via dual pseudo-label supervision." 2022 IEEE International Conference on Image Processing (ICIP). IEEE, 2022.
 - [22] Wang, Z and *et al.* "Weakly-Supervised Self-Ensembling Vision Transformer for MRI Cardiac Segmentation." 2023 IEEE Conference on Artificial Intelligence. IEEE, 2023.
 - [23] Delbrouck, Jean-Benoit, *et al.* "A transformer-based joint-encoding for emotion recognition and sentiment analysis." arXiv preprint arXiv:2006.15955, 2020.
 - [24] Ye, Zhe, Fang Li, and Timothy Baldwin. "Encoding sentiment information into word vectors for sentiment analysis." Proceedings of the 27th international conference on computational linguistics. 2018.
 - [25] Vig, Jesse. "A multiscale visualization of attention in the transformer model." arXiv preprint arXiv:1906.05714, 2019.
 - [26] Wang, and *et al.* "RAR-U-Net: a residual encoder to attention decoder by residual connections framework for spine segmentation under noisy labels." 2021 IEEE International Conference on Image Processing. IEEE, 2021.
 - [27] Kandhro, Irfan Ali, *et al.* "Performance analysis of hyperparameters on a sentiment analysis model." Engineering, Technology & Applied Science Research 10.4, 2020, 6016-6020.
 - [28] Xu, Guixian, *et al.* "Sentiment analysis of comment texts based on BiLSTM." Ieee Access 7, 2019, 51522-51532.
 - [29] Sikarwar, Sandeep Singh, and Nirupma Tiwari. "Analysis the sentiments of amazon reviews dataset by using linear Svc and voting classifier." International journal of science and technology research 9.6, 2020, 461-465.
 - [30] Elbagir, S., & Yang, J. 2018, December. Sentiment analysis of twitter data using machine learning techniques and scikit-learn. In Proceedings of the 2018 International Conference on Algorithms, Computing and Artificial Intelligence (pp. 1-5).
 - [31] Jemai, Fatma, Mohamed Hayouni, and Sahbi Baccar. "Sentiment Analysis Using Machine Learning Algorithms." 2021 International Wireless Communications and Mobile Computing. IEEE, 2021.
 - [32] Bahrainian, Seyed-Ali, and Andreas Dengel. "Sentiment analysis using sentiment features." 2013 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT). Vol. 3. IEEE, 2013.