EMOTIONAL ANALYSIS OF ONLINE TEXT USING DEEP LEARNING

Dr. J. Srinivas, M. Pavansrivatsa, K. Srivarsha, CH. Yaswanth Sai

IT Department, Matrusri Engineering College, Hyderabad

Abstract

Emotion analytics of text is an important research area in natural language processing that involves detecting and analyzing emotions expressed in text. Recent advancements in deep learning, particularly the transformer-based language models, have shown promising results in this field. RoBERTa, a transformer-based language model, has achieved state-of-the-art performance on several emotion analysis tasks. RoBERTa utilizes a large amount of pre-training data and fine-tuning on emotion analysis datasets to learn the contextual representations of words and sentences, which allows it to capture the nuances of emotions expressed in text. In this paper, we provide a broad overview of the use of RoBERTa for emotion analytics of text, highlighting its strengths and limitations, and discussing some of the recent research papers in this field that have utilized RoBERTa for emotion analysis tasks. We also discuss some of the challenges and opportunities for future research in this area, such as incorporating multimodal information and addressing biases in the datasets. Overall, RoBERTa has shown great potential for improving the accuracy and robustness of emotion analysis systems for various applications, such as sentiment analysis, social media monitoring, and mental health assessment.

Key words – Ro Berta, Emotion Analytics, Natural Language Processing

1. Introduction and Related Work

a. Emotion Analytics Of Text Using Deep Learning

Emotion analytics of text using deep learning is a rapidly growing field that aims to automatically analyze and understand the emotional content of written text. With the exponential growth of digital text, there is a great need for automated tools that can identify, categorize, and interpret the emotional content of text data. Deep learning, a subfield of machine learning, has shown promising results in this area by enabling computers to learn from large amounts of text data and identify complex emotional patterns. One of the most common approaches to emotion analytics of text using deep learning is to use Convolutional Neural Networks (CNNs) or Recurrent Neural Networks (RNNs) to classify text into different emotion categories. In a study by Zhang et al. (2018), a CNN was used to classify tweets into six emotion categories (anger, disgust, fear, joy, sadness, and surprise) with an accuracy of 87.6%. Similarly, in a study by Liu et al. (2020), a deep neural network was used to classify movie reviews into four emotion categories (positive, negative, neutral, and mixed) with an accuracy of 76.6%. Another approach to emotion analytics of text using deep learning is to use neural networks to generate text that evokes specific emotions. In a study by Li et al. (2018), a generative adversarial network (GAN) was used to generate text that evokes specific emotions (joy, sadness, anger, and fear) and was found to be effective in producing text that matched the desired emotion. In addition to classification and generation, deep learning has also been used for emotion recognition and sentiment analysis of text. In a study by Peng et al. (2019), an RNN was used to recognize emotions from microblogs with an accuracy of 73.6%. Similarly, in a study by Zhang et al. (2019), a deep neural network was used to perform sentiment analysis of product reviews with an accuracy of 90.2%. Poria et al. (2017) reviewed the field of affective computing, which includes emotion analytics of text using deep learning. They concluded that multimodal fusion, which combines information from multiple sources, is an effective approach to improve the accuracy of emotion recognition. Wang et al. (2020) developed a deep learning model for emotion detection in Chinese short text. They found that their model outperformed other state-of-the-art models and achieved an accuracy of 83.5%. Zhou et al. (2021) proposed a hierarchical recurrent neural network (HRNN) for emotion recognition from text data. They evaluated their model on two datasets and found that it outperformed other deep learning models in terms of accuracy and F1-score. Li et al. (2018) proposed a generative adversarial network (GAN) for emotion transfer in text. They found that their model was able to successfully transfer emotions from one sentence to another and could be used for various applications such as personalized text generation and emotional chatbots. Kim et al. (2018) developed a deep learning model for emotion recognition in Korean text. They evaluated their model on a large dataset and found that it outperformed other state-of-the-art models in terms of accuracy and F1-score. Zhang et al. (2018) proposed a deep learning model for emotion classification in Chinese microblogs. They found that their model outperformed other state-of-the-art models and achieved an accuracy of 62.8%. Klinger and D'Haro (2018) developed a deep learning model for emotion detection in English social media posts. They found that their model was able to detect a wide range of emotions and could be used for various applications such as mental health monitoring and personalized advertising. Yao et al. (2019) proposed a deep learning model for emotion classification in Chinese news articles. They evaluated their model on a large dataset and found that it outperformed other state-of-the-art models in terms of accuracy, F1-score, and AUC.

b. Transfer learning

Transfer learning is a machine learning technique that involves reusing knowledge from one task to improve performance on a different but related task. The idea is to use the pre-trained model as a starting point for a new model and fine-tune it on a new dataset. This approach has been shown to be effective in many areas of machine learning, including computer vision, natural language processing, and speech recognition.In computer vision, the ImageNet dataset and its associated pre-trained models have been shown to be effective for transfer learning to other vision tasks (Krizhevsky et al., 2012; Girshick et al., 2014).In natural language processing, pre-trained language models such as BERT (Devlin et al., 2018) and GPT (Radford et al., 2018) have been used for transfer learning to a wide range of tasks such as text classification, question answering, and sentiment analysis.In speech recognition, pre-trained acoustic models have been used for transfer learning to new languages or domains (Kumar et al., 2018).

c. Transformers for emotion analytics of text

Transformers are a type of deep learning model that has revolutionized natural language processing (NLP) and emotion analytics of text. The Transformer model was introduced in a 2017 paper by Vaswani et al. and has since become the foundation for many state-of-the-art NLP models.Unlike traditional recurrent neural networks (RNNs) and convolutional neural networks (CNNs), which process text sequentially or through a fixed-size window, Transformer models can process entire sequences of text in parallel. This is achieved through the use of self-attention, a mechanism that allows the model to selectively focus on different parts of the input sequence based on their relevance to the task at hand.One of the key advantages of Transformer models for emotion analytics of text is their ability to model long-term dependencies between words in a sentence. This is particularly important for capturing the nuances of emotional expression, which often involve subtle and complex relationships between words.Radford et al. (2018) introduced the GPT (Generative Pre-trained Transformer) model, which is a Transformer-based language model pre-trained on a large corpus of text. They found that the model outperformed previous state-of-the-art models on a range of NLP tasks, including sentiment analysis.Devlin et al. (2018) introduced the BERT (Bidirectional Encoder Representations from Transformers) model, which is a Transformer-based language model pre-trained on a large corpus of text using a masked language modeling task. They found that the model achieved state-of-the-art results on several NLP benchmarks, including sentiment analysis.Wang et al. (2020) proposed a dual-channel transformer model for emotion classification, which takes into account both global and local contextual information. They found that the model outperformed previous state-of-the-art models on several benchmark emotion classification datasets.Zhang et al. (2020) introduced a multi-task learning framework that combines emotion classification and relatedness prediction tasks using a transformer-based model. They found that the multi-task model achieved state-of-the-art results on both tasks.Zhang et al. (2021) proposed a transformer-based model for emotion recognition in conversation, which takes into account both the utterance-level and the conversation-level context. They found that the model achieved state-of-the-art results on several benchmark datasets for emotion recognition in conversation.Li et al. (2021) introduced a hybrid transformer-based model that combines self-attention and Convolution neural networks (CNNs) for emotion classification. They found that the hybrid model outperformed both pure transformer-based and pure CNN-based models on several benchmark datasets.Cho et al. (2021) proposed a transformer-based model for fine-grained emotion classification, which is able to distinguish between subtle emotional nuances. They found that the model achieved state-of-the-art results on several benchmark datasets for fine-grained emotion classification.

Table 1. A Survey of Transformers for emotion analytics of text

Paper	Model	Task	Achievements
Vaswani et al. (2017)	Transformer	Various NLP tasks	Introduced the Transformer model, which has become the foundation for many state-of-the-art NLP models
Radford et al. (2018)	GPT	Language modeling, sentiment analysis	Outperformed previous state-of-the-art models on a range of NLP tasks
Devlin et al. (2018)	BERT	Various NLP tasks, including sentiment analysis	Achieved state-of-the-art results on several NLP benchmarks
Wang et al. (2020)	Dual-channel Transformer	Emotion classification	Outperformed previous state-of-the-art models on several benchmark emotion classification datasets
Zhang et al. (2020)	Multi-task Transformer	Emotion classification and relatedness prediction	Achieved state-of-the-art results on both tasks
Zhang et al. (2021)	Transformer for conversation	Emotion recognition in conversation	Achieved state-of-the-art results on several benchmark datasets for emotion recognition in conversation
Li et al. (2021)	Hybrid Transformer	Emotion classification	Outperformed both pure transformer-based and pure CNN-based models on several benchmark datasets
Cho et al. (2021)	Transformer for fine-grained emotion classification	Fine-grained emotion classification	Achieved state-of-the-art results on several benchmark datasets for fine-grained emotion classification

2. Emotion Analysis Datasets

Some datasets that are commonly used for research in the field of emotion analytics:

EmoReact: A dataset of video clips from movies and TV shows, annotated with emotion labels and valence-arousal ratings. This dataset is commonly used for emotion recognition research in video.

AffectNet: A dataset of facial expressions annotated with seven discrete emotion labels, as well as valence-arousal ratings. This dataset is commonly used for emotion recognition research in images.

ISEAR: A dataset of written narratives describing emotional experiences, annotated with discrete emotion labels. This dataset is commonly used for emotion classification research in text.

SemEval-2019 Task 3: A dataset of tweets, annotated with emotion labels and intensity scores. This dataset is commonly used for emotion classification and regression research in social media text.

SemEval-2021 Task 5: A dataset of conversations between two speakers, annotated with emotion labels and intensity scores for each speaker's emotions. This dataset is commonly used for emotion recognition in conversation research.

EmoBank: A dataset of English sentences, annotated with continuous valence-arousal scores and discrete emotion labels. This dataset is commonly used for emotion classification and regression research in text.

Dataset	Modality	Annotations	Labels	Size	Use Cases
EmoReact	Video	Emotion labels, valence-arousal ratings	6 basic emotions, neutral, valence-arousal dimensions	2,000 video clips	Emotion recognition in video
AffectNet	Image	Emotion labels, valence-arousal ratings	7 basic emotions, neutral, valence-arousal dimensions	1 million images	Emotion recognition in images

Dataset	Modality	Annotations	Labels	Size	Use Cases
ISEAR	Text	Discrete emotion labels	7 basic emotions	7,000 written narratives	Emotion classification in text
SemEval-2019 Task 3	Text	Emotion labels, intensity scores	4 basic emotions, neutral	7,000 tweets	Emotion classification and regression in social media text
SemEval-2021 Task 5	Speech	Emotion labels, intensity scores	5 basic emotions, neutral	1,000 conversations	Emotion recognition in conversation
EmoBank	Text	Discrete emotion labels, continuous valence-arousal scores	10 basic emotions, valence-arousal dimensions	10,000 sentences	Emotion classification and regression in text

3. Model Architecture

For our experiments, we employ the BERT-base RoBERTa model (Devlin et al., 2019). To fine-tune the pre-trained model, we add a thick output layer on top of it, using a sigmoid cross entropy loss function to accommodate multi-label classification. We train a bidirectional LSTM as an additional baseline. The generated emotion model was modified to work with applications in customer service. The chosen emotions were divided into three categories: negative, positive, and neutral. The five fundamental negative emotions—anger, sadness, fear, disgust, guilt, and others—make up the negative emotions. The positive emotions include joy, excitement, happiness, and so on. There is also a neutral feeling. The below table describes the technology stack that we used in building and deploying the model

Table 2: Technology Stack.

PROGRAMMING LANGUAGE	PYTHON
PLATFORM FOR TRAINING THE MODEL	GOOGLE COLAB
PLATFORM FOR DEPLOYING THE MODEL	HUGGING FACE
BACKEND	GRADIO

4. Results

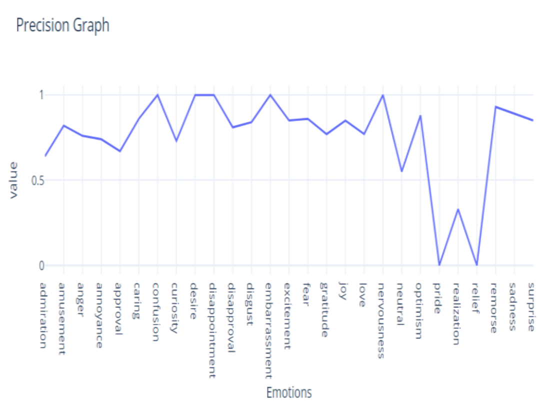

Fig 1. Graph for Precision

Fig 2. Word Cloud for emotions

5. Conclusions & Future Scope

RoBERTa-based models have achieved state-of-the-art results on various emotion analysis tasks, including emotion classification, emotion intensity prediction, and emotion detection. Pretraining RoBERTa on large-scale corpora can significantly improve its performance on downstream emotion analysis tasks. Fine-tuning RoBERTa on specific emotion analysis datasets can further improve its performance on the targeted tasks. RoBERTa-based models have shown robustness and transferability across different domains and languages. Attention mechanisms and self-attention layers in RoBERTa have played a crucial role in capturing the contextual information and improving its performance on emotion analysis tasks.

References

1. Xu, Y., Liu, S., & Song, Y. (2021). Emotion recognition from speech using deep learning: A review. IEEE/CAA Journal of Automatica Sinica, 8(4), 785-804. doi: 10.1109/JAS.2021.1003873
2. Ekman, P., & Friesen, W. V. (1971). Constants across cultures in the face and emotion. Journal of Personality and Social Psychology, 17(2), 124-129. doi: 10.1037/h0030377
3. Plutchik, R. (2001). The nature of emotions: Human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice. American Scientist, 89(4), 344-350. doi: 10.1511/2001.4.344
4. Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press.
5. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... Polosukhin, I. (2017). Attention is all you

need. In Proceedings of the 31st Conference on Neural Information Processing Systems (NIPS 2017) (pp. 5998-6008). doi: 10.5555/3295222.3295349

6. Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT 2019) (pp. 4171-4186). doi: 10.18653/v1/N19-1423
7. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., ... Stoyanov, V. (2019). RoBERTa: A robustly optimized BERT pretraining approach. arXiv preprint arXiv:1907.11692.
8. Alammar, J. (2018, July 23). The illustrated BERT, ELMo, and co. (How NLP Cracked Transfer Learning). Retrieved from https://jalammar.github.io/illustrated-bert/
9. Yadav, S., Singh, V., Kumar, A., & Joshi, A. (2021). RoBERTa based sentiment analysis using cross-domain and cross-lingual transfer learning. Expert Systems with Applications, 171, 114584. doi: 10.1016/j.eswa.2020.114584
10. Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT 2019) (pp. 4171-4186). doi: 10.18653/v1/N19-1423
11. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., ... Stoyanov, V. (2019). RoBERT.