

# Optimizing Small Transformer-Based Language Models for Multi-Label Sentiment Analysis in Short Texts

Julius Neumann<sup>1\*</sup>, Robert Lange<sup>1\*</sup>, Yuni Susanti<sup>2</sup>, and Michael Färber<sup>1</sup>

<sup>1</sup> ScaDS.AI, TU Dresden, Germany

[michael.faerber@tu-dresden.de](mailto:michael.faerber@tu-dresden.de)

<sup>2</sup> FIZ Karlsruhe, Germany

**Abstract.** Sentiment classification for short texts faces significant challenges such as class imbalance, limited training samples, and the inherent subjectivity of sentiment labels – issues that are further intensified by the limited context in short texts. These factors make it difficult to resolve ambiguity and exacerbate data sparsity, hindering effective learning. In this paper, we evaluate the effectiveness of *small* Transformer-based models (i.e., BERT and RoBERTa, with fewer than 1 billion parameters) for multi-label sentiment classification, with a particular focus on short-text settings. Specifically, we evaluated three key factors influencing model performance: (1) continued domain-specific pre-training, (2) data augmentation using automatically generated examples, specifically generative data augmentation, and (3) architectural variations of the classification head. Our experimental results show that data augmentation improves classification performance, while continued pre-training on augmented datasets can introduce noise rather than boost accuracy. Furthermore, we confirm that modifications to the classification head yield only marginal benefits. These findings provide practical guidance for optimizing BERT-based models in resource-constrained settings and refining strategies for sentiment classification in short-text datasets.

**Keywords:** sentiment analysis · short-text datasets · small language models.

## 1 Introduction

In modern practical decision-making, building predictive models often involves significant challenges, including imbalanced class distributions, highly variable dataset sizes – from limited to extremely large – and inherent subjectivity in the ground truth labels. These factors can severely degrade model performance and further complicate learning. Consequently, developing novel learning methods capable of addressing these difficulties is a critical area of ongoing research. Within this context, multi-label sentiment classification in short texts represents

---

\* Equal contribution.

a particularly challenging task, requiring models to infer subjective emotional states (e.g., anger, joy, or surprise) from sparse and ambiguous linguistic cues. Short texts inherently limit contextual information, reducing the availability of disambiguating semantic cues and amplifying the effects of data sparsity. An additional common issue in this task includes the inherently subjective nature of sentiment annotations, which often results in low inter-annotator agreement. Short texts, especially from social media, frequently contain informal language, abbreviations, emojis, and sarcasm, further complicating accurate sentiment detection. Addressing these difficulties is essential for advancing sentiment analysis and related Natural Language Processing (NLP) tasks and applications.

Although pre-trained Transformer-based models such as BERT [4] have significantly advanced sentiment analysis [1,4,8], critical challenges remain, including handling class imbalance, data scarcity, and subtle semantic cues within short texts. Recent research has shown that expanding training data through generative data augmentation [2], refining model architectures [14], and using domain-specific pre-training [7] potentially enhance the model’s performance. However, only a few studies systematically evaluate these approaches in combination. Existing research predominantly focuses on binary sentiment classification or assumes access to large, well-curated datasets, leaving open questions about the optimal strategies for multi-label emotion detection in small, noisy and inherently subjective text corpora.

In this paper, we address these challenges through a comprehensive evaluation of the following three key strategies: (1) continued domain-specific pre-training to better capture linguistic nuances, (2) generative data augmentation (GDA) to alleviate data scarcity, and (3) classification head variations to enhance parameter efficiency and adaptability. Our experiments were carried out on 2,768 short English texts from the SemEval 2025 Task 11 corpus [3], utilizing *small* Transformer-based language models, specifically BERT [4] and its variant RoBERTa [8]. In this work, we define *small* as language models with fewer than 1 billion parameters, following [17]. Such models offer a favorable trade-off between performance and computational efficiency, making them well-suited for deployment in resource-constrained real-world scenarios. Additionally, we apply SHAP [9] analysis to improve the interpretability of the predictions.

The following summarizes our main contributions:

- We systematically evaluate three strategies – continued domain-specific pre-training, generative data augmentation, and classification head variations – for multi-label sentiment classification in short texts, targeting challenges of linguistic sparsity, data scarcity, and model efficiency. We release the code and dataset on GitHub.<sup>3</sup>
- Our results show that moderate generative data augmentation improves BERT-based model performance, while excessive domain-specific pre-training can add noise. Variations in classification heads yield minimal gains. These findings provide practical guidance for BERT optimization in real-world, short-text datasets.

---

<sup>3</sup> <https://github.com/faerber-lab/shorttext-sentiment-transformers>

- We employ SHAP analysis to interpret model predictions, offering insights into feature contributions and enhancing model transparency in short-text sentiment classification tasks.

## 2 Related Work

*Transformer-Based Text Classification.* Transformer-based models such as BERT [4] have become the state of the art for text classification, largely due to their ability to encode contextualized representations with bidirectional attention [1]. Compared to earlier architectures like LSTMs, Transformers more effectively capture long-range semantic dependencies. Beyond model architecture, the quantity of available training data strongly influences performance. [7] and [10] show that larger training sets generally lead to improved generalization in NLP tasks. BERT [4] and its refined variant RoBERTa [8] are often the starting points for a wide range of downstream applications. RoBERTa’s removal of next sentence prediction and its employment of extended pre-training with dynamic masking enhance performance on multiple benchmarks [8,6]. In typical classification scenarios, the hidden representation of the [CLS] token is passed through a classifier, and both the Transformer layers and the classifier head are fine-tuned to maximize the probability of the correct label. Adjustments in the head’s architecture, such as varying embedding size or adopting parameter-efficient layers like projected attention [14], can further enhance performance [20].

*Data Augmentation.* Data augmentation has been widely adopted to improve model generalization in NLP tasks. It typically involves selecting or creating meaningful samples, such as work by [16] or adding extra information [5]. While adversarial approaches have been explored [5,12], their complexity and mixed empirical results have prompted interest in simpler, generative methods. [2] proposes a more straightforward strategy, using GPT-3 to generate new training examples and filtering them to retain only the most promising candidates. Even with a dataset of only 26 examples and two labels, they report substantial accuracy improvements, but also observe that excessive augmentation can introduce noise and ultimately degrade performance [11].

*Generative Data Augmentation.* Although generative data augmentation (GDA) has been explored in previous work [13], we adapt it here for multi-label sentiment classification on short texts. In contrast to Unsupervised Data Augmentation (UDA) [21], which emphasizes consistency training on unlabeled data, we directly generate new examples using a state-of-the-art language model (e.g., GPT-4o) and incorporate them into the training set without additional filtering. We refer to this approach as *Generative Data Augmentation (GDA)*. By verifying that the synthetic samples remain domain-consistent, we aim to capture rare sentiment cues and edge cases often missing from the original dataset. Similar to prior studies, we hypothesize that augmenting too aggressively may introduce noise and lead to performance plateaus or overfitting, but a moderate infusion

of diverse synthetic data can substantively boost our multi-label classification performance.

### 3 Methodology

Our work focuses on investigating the performance of Transformer-based *small* language models for the short-text multi-label sentiment analysis task. Given a short text as the input, the model predicts one or more sentiment labels that best describe the emotional content of the input text. Formally, let  $\mathcal{D} = \{(x_i, Y_i)\}_{i=1}^N$  represent the dataset, where each sample consists of:

- $x_i \in \mathcal{X}$ : a short text sample as the input
- $Y_i \subseteq \mathcal{Y}$ : a subset of sentiment labels assigned to  $x_i$ , where:  
 $\mathcal{Y} = \{\text{Anger, Fear, Joy, Sadness, Surprise}\}$  is the set of all possible sentiment classes.

Unlike single-label classification, each instance in the dataset can be associated with multiple sentiment labels simultaneously. To evaluate model performance, we employ standard multi-label classification metrics following [3], including *Accuracy*, *F1 Score*, *Macro F1 Score*, and *Cohen’s Kappa Score*. The Macro F1 Score ensures a balanced assessment across all sentiment categories, addressing potential class imbalance issues.

#### 3.1 Approach

In this work, we investigate three distinct approaches to enhance the performance of Transformer-based language models, described in detail below. Our experiments focus on BERT and its variant, RoBERTa, as the base models.

**Continued Domain-Specific Pre-training** This approach involves continuing the pre-training of base language models on domain-specific corpora, as recommended by [15]. This strategy is particularly relevant for short-text sentiment classification, where limited context and sparse linguistic cues make capturing domain-specific language nuances critical. Supporting this, [11] demonstrated that domain-focused pre-training leads to significantly greater performance gains than merely increasing the size of the general training dataset beyond a certain threshold, highlighting the importance of specialized knowledge for improving accuracy in contextually constrained scenarios such as short texts.

To perform continued pre-training of the models in our experiment, only the *text* field of the training dataset is utilized, following the work by [15]. The text samples are tokenized using a pre-trained tokenizer, with each model employing its corresponding tokenizer from `HuggingFace` [19] library. The models are then trained on these tokenized text samples using the masked language modeling (MLM) objective, wherein tokens are randomly masked with a probability of 15%. Pre-training is carried out for 50 epochs with a learning rate of  $2 \times 10^{-5}$  and a weight decay of 0.01.

**Generative Data Augmentation (GDA)** This approach enhances the original dataset by incorporating automatically generated samples from a generative model, which is particularly beneficial for our task where data scarcity poses significant challenges. By expanding the training data with synthetic examples that reflect the original dataset-specific language patterns, GDA helps models better capture subtle linguistic nuances that are often underrepresented in particularly small datasets. Prior studies have demonstrated that such augmentation can significantly improve model performance in certain cases [7,10].

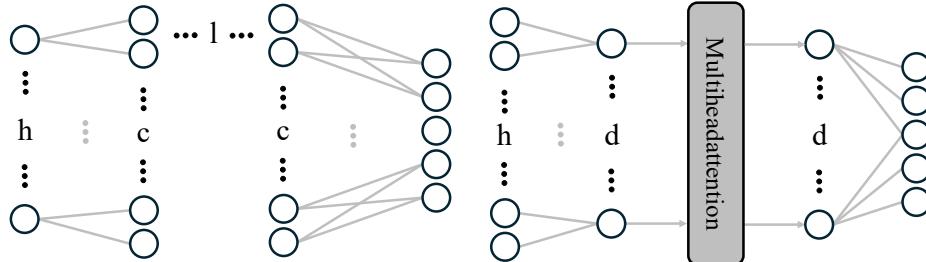
To augment the original training set in our experiment, we utilized the OpenAI completions API endpoint with a fine-tuned `GPT-4o-mini model`. This model was fine-tuned for three epochs on a partitioned training set incorporating sentiment labels, to match the label distribution of the dataset under test. The system prompt specified the following instruction:

Generate short texts with their corresponding sentiment labels. The sentiment labels include Anger, Fear, Joy, Sadness, and Surprise. The texts are in English and have a maximum length of 256 characters.

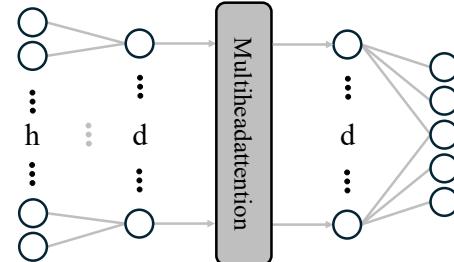
For each generated example, we provided the same prompt along with an output template in a `JSON` schema, which included a text field and an array of assigned sentiment labels. This procedure was repeated until a total of 11,684 augmented examples were generated, a number chosen to align with our available computational resources. Since the generated texts exhibited slight variations in average length compared to the original dataset, many samples were split into two or three parts to better align with the original training domain. Subsequent fine-tuning on the augmented dataset resulted in improved model performance, reducing label ambiguity, and confirming that the generated data maintained a quality comparable to that of the original dataset. The class distribution of the original dataset and the dataset after augmentation using the generative data augmentation method are presented in Tables 1 and 2, respectively.

**Classification Head Architectures** In the third approach, we evaluate different architectures for the model’s classification head, a component especially critical for short-text sentiment classification where limited input length demands efficient extraction and integration of the model features. As highlighted by [18], the design of the feed-forward classification head can substantially impact overall model performance. We compare two distinct classification head architectures: (1) a *fully-connected layer* and (2) *projected attention*, described in detail below.

A *fully-connected layer* architecture provides a straightforward mapping from encoded features to labels. Concretely, it consists of a fully-connected layer network with intermediate `ReLU` activation function and dropout layers, as illustrated in Figure 1. In this setting, the hidden size  $h$  of the pre-trained model serves as the input dimension. This is followed by  $l$  linear layers, each with an internal classifier size  $c$ . The final classification layer produces an output of size



**Fig. 1.** Illustration of the *fully-connected layer* classification head.



**Fig. 2.** Illustration of the *projected attention* classification head.

equal to the number of labels. Finally, a *sigmoid* function is applied to compute the probability for each label, following the practice in multi-label classification.

The second architecture employs the *projected attention* mechanism, as proposed by [14]. This technique incorporates an attention mechanism (see Figure 2) to dynamically focus on important features – potentially enhancing the model’s ability to capture nuanced sentiment signals in sparse, context-limited short texts. In this architecture, the classification head consists of a linear projection layer that transforms the pre-trained model’s output size  $h$  to the attention dimension  $d$ . This is followed by a multi-head attention mechanism with  $k$  attention heads and an embedding dimension of  $d$ . Lastly, an attention-based classifier projects the attention output to match the number of labels.

### 3.2 Dataset

Our experiments are conducted on the SemEval 2025 Task 11 dataset [3], which comprises 2,768 short English texts, each labeled with zero, one, or multiple sentiment classes: {Anger, Fear, Joy, Sadness, and Surprise}. On average, each text consists of 78.4 characters, corresponding to approximately 15 words, which categorizes this task as a “short text” classification task.

The distribution of labeled classes in the training set is presented in Table 1. The dataset exhibits a class imbalance, with the “Fear” class being over-represented (58.2% of the total instances) and the “Anger” class being under-represented (12.0% of the total instances). This may introduce bias in the classification model, potentially leading to a higher propensity for predicting the over-represented class (“Fear” class) and a reduced likelihood of correctly identifying instances of the under-represented class (“Anger” class). The class distribution of the dataset after augmentation with the generative data augmentation method is presented in Table 2.

### 3.3 Experimental Details

In the experiment, we consider four different variants of BERT and RoBERTa pre-trained models as base models for the task-specific fine-tuning: (1) `bert-base`,

**Table 1.** Original dataset class distribution.

Label	Frequency	Probability (%)
Anger	333	12.0
Fear	1611	58.2
Joy	674	24.3
Sadness	878	31.7
Surprise	839	30.3

**Table 2.** Augmented dataset class distribution.

Label	Frequency	Probability (%)
Anger	2227	19.0
Fear	7606	65.0
Joy	3577	30.6
Sadness	4884	41.8
Surprise	4645	39.7

(2) `bert-large`, (3) `roberta-base`, and (4) `roberta-large`. Each model is initialized from a checkpoint available in `HuggingFace` [19]. This setup enables a comparative analysis of BERT and RoBERTa model architectures while also assessing the impact of model size.

In each training configuration, all four models are fine-tuned twice: once without continued pre-training and once with additional pre-training on domain-specific data. In the first configuration, the models are pre-trained and fine-tuned on the original training dataset provided by [3]. In the second configuration, 33.3% of the artificially generated samples are incorporated to expand the training dataset. This process is then repeated twice, once with 66.6% of the generated samples and again with 100% of the generated samples, to analyze whether performance improves further or diminishes. Before tokenization, all words in the training dataset are converted to lowercase, and hyperlinks as well as punctuation marks are removed.

For all configurations, the classification head consists of a fully connected network with  $l = 1$  linear layer and an internal classifier size of  $c = 768$ , which matches the hidden size of the BERT models. To determine whether a different classification head size or architecture improves model performance, the best-performing base model from the initial trials is further fine-tuned with a fully connected classification head containing  $l = 2$  and  $l = 4$  linear layers. Additionally, fine-tuning is performed with a “Projected Attention” head incorporating  $k = 1$  and  $k = 2$  attention heads.

All models are fine-tuned for 30 epochs with a learning rate of  $2 \times 10^{-5}$ , a weight decay of 0.01, `AdamW` optimizer, and a training and evaluation batch size of 32. All experiments were conducted on NVIDIA A100 and V100 GPUs, using the single, predefined train-validation-test split specified by the original dataset to ensure consistency with prior works. Other hyperparameter settings and implementation details are provided in our GitHub.

### 3.4 Explainability and Reasoning Metrics

To further understand the reasoning behind model predictions, we generate confusion matrices and employ explainability library `SHAP` (SHapley Additive exPlanations) [9] to analyze feature importance in classification decisions. Confusion matrix analysis helps identify common misclassifications and potential biases,

**Table 3.** Model performance across different training variations.

Model	Pretrained	Original Dataset		0.33 Augmentation		0.66 Augmentation		1.0 Augmentation	
		Acc.	F1 Score	Acc.	F1 Score	Acc.	F1 Score	Acc.	F1 Score
<code>bert-base</code>	Yes	0.365	0.694	0.375	0.687	0.365	0.71	0.438	0.706
	No	0.344	0.677	0.313	0.686	0.406	0.715	0.354	0.666
<code>bert-large</code>	Yes	0.427	0.697	0.406	0.708	0.375	0.697	-	-
	No	0.354	0.689	0.365	0.698	0.385	0.703	0.406	0.715
<code>roberta-base</code>	Yes	0.417	0.707	0.365	0.678	0.406	0.693	0.427	0.719
	No	0.396	0.711	0.396	0.698	0.406	0.706	0.458	0.749
<code>roberta-large</code>	Yes	0.448	0.73	0.406	0.74	0.417	0.736	0.396	0.733
	No	<b>0.469</b>	<b>0.77</b>	0.396	0.719	<b>0.469</b>	0.764	0.458	0.767

providing valuable information about the limitations of the models’ decision-making process. Meanwhile, SHAP utilizes *Shapley* values to compute feature contribution scores, enabling a detailed assessment of input feature importance. By applying SHAP analysis, we can determine the impact of specific words on sentiment classification, offering a more interpretable understanding of model behavior.

## 4 Evaluation

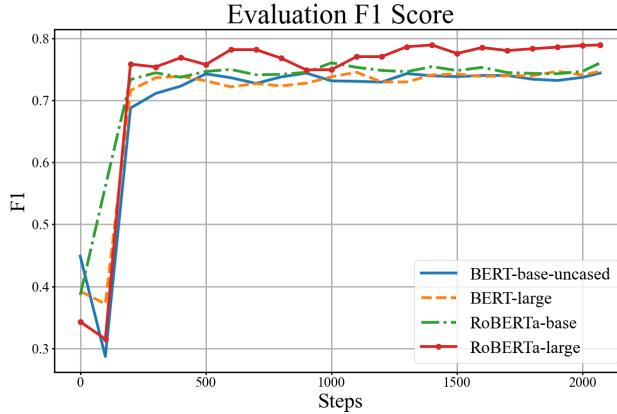
We conducted evaluation experiments to assess the impact of the three strategies outlined in Section 3.1, *Approach*. Additionally, we performed an explainability analysis to gain deeper insights into the model’s decision-making process. The main experimental findings and explainability results are presented in the following sections.

### 4.1 Main Results

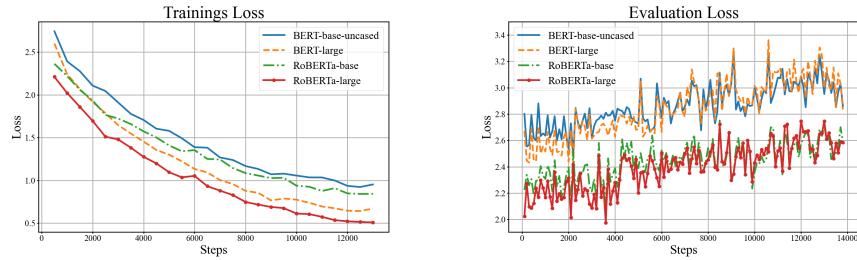
The evaluation results for continued pre-training of the models and data augmentation techniques are summarized in Table 3 and illustrated in Figure 3, while the experiment results for classification head architecture variation are presented in Table 4.

As shown in Figure 3, after approximately 500 training steps, the F1 Score of the `bert-base`, `bert-large`, and `roberta-base` models plateaus, indicating minimal further improvement. In contrast, the `roberta-large` model continues to benefit from extended fine-tuning. Additionally, while larger models generally achieve better performance, the performance gap between `bert-base` and `bert-large` on the unmodified training set is marginal, with `bert-base` even outperforming `bert-large` in one instance. Conversely, the `roberta-large` model consistently outperforms `roberta-base` and is generally the best-performing model overall.

Overall, although larger models, such as `bert-large` and `roberta-large`, generally perform better, smaller models can be improved effectively through data augmentation and domain-specific pre-training. Among these, data augmentation proved most effective, improving all models except `roberta-large`. Modifying the classification head had little impact on performance but may reduce



**Fig. 3.** F1 Scores of the four base models with continued pre-training and fine-tuning on the original dataset.



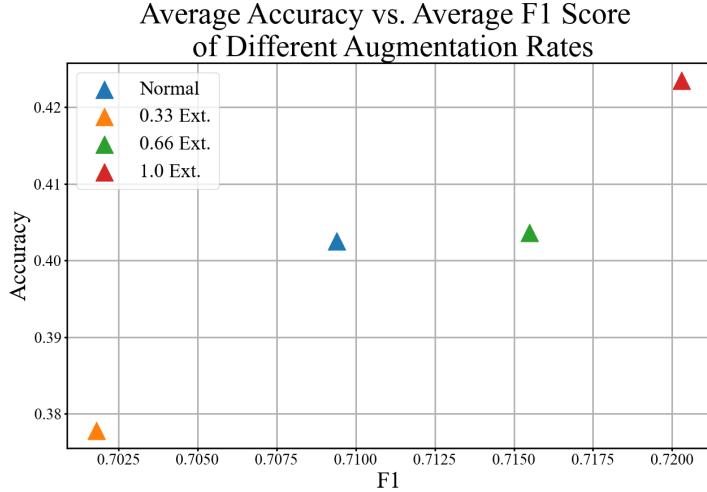
**Fig. 4.** Training loss of the four base models during continued pre-training on the original dataset.

**Fig. 5.** Evaluation loss of the four base models during continued pre-training on the original dataset.

trainable parameters, which is beneficial in resource-constrained settings. In the following, we discuss the impact of each of the three approaches more in detail.

*Impact of the Continued Domain-Specific Pre-Training.* When comparing the evaluation loss (Figure 5) with the training loss (Figure 4) during pre-training, it becomes evident that overfitting occurs rapidly, leading to diminishing returns. Specifically, continued domain-specific pre-training appeared to slightly improve the performance of the `bert-base` and `bert-large` models in most cases. However, the `roberta-base` and `roberta-large` models generally performed better without additional pre-training. This may be attributed to the fact that RoBERTa models were pre-trained for a longer duration on larger corpus, making continued pre-training prone to overfitting and a subsequent loss of generalization.

*Impact of Generative Data Augmentation.* The evaluation results of the model for different data augmentation rates are shown in Figure 6 and Table 3. Overall,



**Fig. 6.** Average F1 Scores and accuracy across different data augmentation rates.

accuracy and F1 scores increase as the training set expands, although gains are less than expected. Notably, smaller models such as `bert-base` show bigger performance improvements compared to larger models e.g., `roberta-large`.

A slight expansion of the training set initially results in a small drop in performance relative to the unmodified dataset. However, as the dataset size increases, performance improves: the dataset with 66% augmentation performs comparably to the original, while the fully augmented dataset achieves the highest results – improving accuracy by two percentage points and F1 Score by an average of 0.01. In contrast, the `roberta-large` model consistently experiences performance degradation across all levels of data augmentation.

These findings suggest that the dataset size threshold, as described by [2], has not yet been reached. A further expansion of the training set might yield additional performance gains. Pre-training on the extended dataset appears to be less effective than direct fine-tuning, likely due to the introduction of noise into the domain-specific training data.

Overall, our results demonstrate that generative data augmentation is an effective strategy, which leads to a moderate improvement in model performance. By enriching the training domain with unmoderated examples generated by a state-of-the-art text generation model, we achieved measurable improvements in model performance. These findings underscore the potential of leveraging generative techniques to enhance the diversity and quality of training data.

*Impact of Classification Head Architecture Variation.* The evaluation results of the model with different classification heads is summarized in Table 4. Overall, increasing the number of parameters in the classification head did not yield significant improvements in model performance. In fact, larger classification heads resulted in diminished performance, likely due to overfitting. Despite having fewer

**Table 4.** Evaluation results of `roberta-large` model with different classification heads (fc-*fully-connected* and Proj. Att-*projected attention*).

Class. Head	Accuracy	F1 Score	Params
fc 768x2	<b>0.469</b>	<b>0.77</b>	1x
fc 768x4	0.417	0.743	3x
fc 1536x2	0.458	0.77	4x
fc 1536x4	0.455	0.756	12x
Proj. Att. 128x1	0.432	0.743	<b>0.44x</b>
Proj. Att. 128x2	0.458	0.756	0.61x

parameters, models utilizing the projected attention head demonstrated competitive effectiveness. This suggests that a well-designed attention-based classification head can enhance performance efficiency, achieving better results per parameter compared to standard fully connected layers.

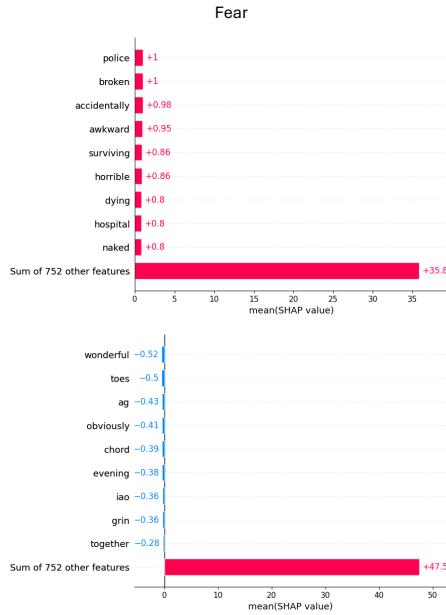
#### 4.2 Explainability Analysis Results

Figure 7 shows the class-specific influence graph for the “Fear” and “Anger” classes produced by SHAP analysis, with the graphs and explanation for other classes provided in the *Appendix*. These visualizations offer insights into the key features influencing classification decisions by the model. For instance, the token *police* emerges as a strong predictor for the “Fear” class, indicating that the model has learned to associate this token with expressions of fear in the dataset. For the “Anger” class, *insulting* and *furious* are a strong positive predictor, while *ailments* and *laughed* contribute minimally. These influence patterns serve as baseline indicators of the model’s interpretability. The SHAP analysis further reveals that the models construct a structured representation of sentiment, leveraging specific words and word combinations to make predictions.

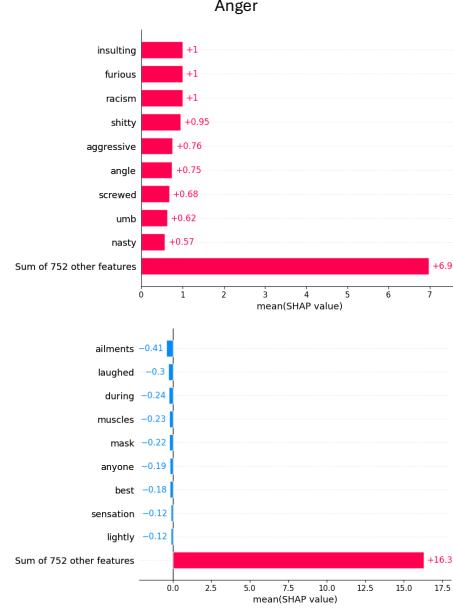
Furthermore, the confusion matrices in Figure 9 offer deeper insights into the challenges of sentiment classification in short texts. They reveal that classes with higher prior probabilities are more accurately classified, while low-frequency classes suffer from reduced accuracy, reflecting the model’s reliance on the underlying data distribution when making decisions. For instance, the high number of false positives for the “Fear” class suggests a learned bias, providing insight into how the model internalizes and applies decision boundaries. This underscores how both the quality of the training data and the complexity of emotional expression influence model interpretability.

#### 4.3 Reviewer Agreement and Human Evaluation

Although an F1 Score below 0.9 and an accuracy under 80% may initially seem suboptimal, it is crucial to recognize that multi-label sentiment analysis is inherently subjective. The high-dimensional nature of sentiment data and the limited



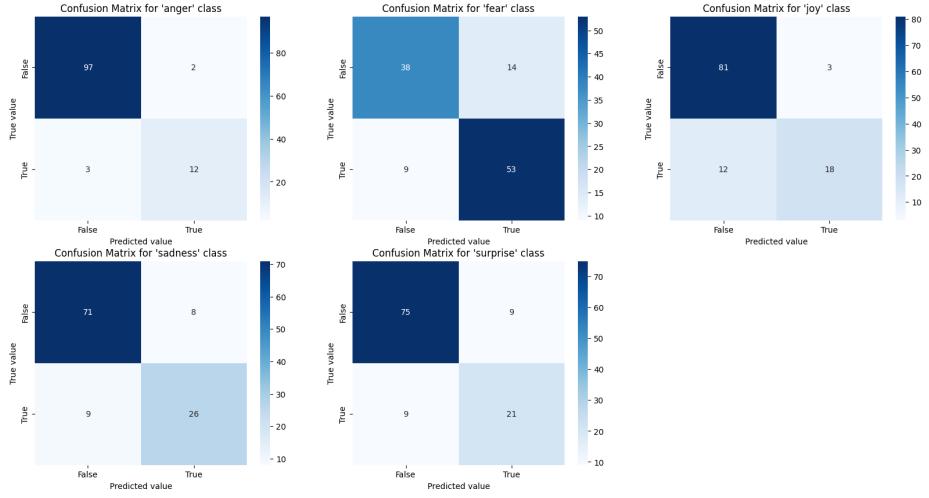
**Fig. 7.** SHAP graph for “Fear” emotion class.



**Fig. 8.** SHAP graph for “Anger” emotion class.

long-context understanding in short texts often prevent exceptionally high accuracy. To further investigate this subjectivity, we conducted a human evaluation in which three English speakers (*non-native*) annotated a subset of texts from the development set and compared it against the gold truth. This assessment provided insights into annotation variability and the inherent challenges of sentiment classification.

The results, as summarized in Table 5 and Table 6, show that our model significantly outperforms human annotators, demonstrating its robustness in handling nuanced sentiment distinctions. Notably, our best model, despite an accuracy of only 0.469 and an F1 Score of 0.77, outperformed human evaluators, underscoring its capability in sentiment classification. Moreover, the moderate Cohen’s Kappa Scores (ranging from 0.293 for “Fear” to 0.709 for “Joy”) highlight the difficulty of achieving consistent sentiment annotations. Lower agreement for emotions such as *Fear*, *Sadness*, and *Surprise* suggests that these categories are inherently more ambiguous, particularly for non-native speakers. This variability further supports our observation that the model effectively navigates the complexities and subjectivity inherent in sentiment classification.

**Fig. 9.** Confusion matrix of the model.**Table 5.** Human evaluations vs. our model.

	Accuracy F1 Score	
Human Evaluator 1	0.224	0.549
Human Evaluator 2	0.290	0.612
Human Evaluator 3	0.252	0.578
Our Model	0.469	0.770

**Table 6.** Cohen's Kappa Score between human evaluators.

	Cohen's Kappa
Anger	0.504
Fear	0.293
Joy	0.709
Sadness	0.307
Surprise	0.346

## 5 Conclusion

In this study, we investigated the impact of three key factors – (1) continued domain-specific pre-training, (2) generative data augmentation (GDA), and (3) classification head architecture – on multi-label sentiment classification in short texts. Given the limited contextual information typical of short-text datasets, these factors are especially critical for improving the performance of *small* Transformer-based language models, e.g., BERT and RoBERTa.

Our results indicate that while larger-parameter models (e.g., `roberta-large`) generally outperform smaller-parameter models (`bert-base` and `roberta-base`), moderate application of GDA consistently enhances the performance of the smaller models, suggesting that synthetic samples effectively mitigate data scarcity. Continued pre-training was beneficial for BERT-based architectures but introduced noise in RoBERTa variants, highlighting the model-specific nature of domain adaptation for short texts. Although modifications to the classification head had minimal impact on accuracy, they may reduce the number of train-

able parameters – an important consideration for resource-constrained settings common in real-world applications.

Our findings advance the development of tailored learning methods to overcome practical data difficulties, offering actionable insights for optimizing multi-label sentiment classification in short texts – a domain characterized by sparse and ambiguous data. Future work could explore refined filtering strategies for GDA and evaluate a broader range of architectures to further improve performance, particularly in resource-limited and complex data scenarios.

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

## References

1. Alaparthi, S., Mishra, M.: Bidirectional encoder representations from transformers (bert): A sentiment analysis odyssey. ArXiv **abs/2007.01127** (2020), <https://api.semanticscholar.org/CorpusID:220302509>
2. Balkus, S.V., Yan, D.: Improving short text classification with augmented data using gpt-3. Natural Language Engineering **30**(5), 943–972 (2024). <https://doi.org/10.1017/S1351324923000438>
3. Codabench: Bridging the gap in text-based emotion detection - semeval 2025 task 11 - track a. Codabench (2024-25), <https://www.codabench.org/competitions/3863>, accessed: 2025-01-13
4. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding (2019), <https://arxiv.org/abs/1810.04805>
5. Guo, X., Yu, H.: On the domain adaptation and generalization of pretrained language models: A survey (2022), <https://arxiv.org/abs/2211.03154>
6. Hirlea, D., Bryant, C., Rei, M.: Contextual sentence classification: Detecting sustainability initiatives in company reports. ArXiv **abs/2110.03727** (2021), <https://api.semanticscholar.org/CorpusID:238531297>
7. Ji, Y., Deng, Y., Gong, Y., Peng, Y., Niu, Q., Zhang, L., Ma, B., Li, X.: Exploring the impact of instruction data scaling on large language models: An empirical study on real-world use cases. ArXiv **abs/2303.14742** (2023), <https://api.semanticscholar.org/CorpusID:257766844>
8. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach (2019), <https://arxiv.org/abs/1907.11692>
9. Lundberg, S.M., Lee, S.I.: Shap. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 30. Curran Associates, Inc. (2017), [https://proceedings.neurips.cc/paper\\_files/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf)
10. Mehrafarin, H., Rajaee, S., Pilehvar, M.T.: On the importance of data size in probing fine-tuned models. In: Findings (2022), <https://api.semanticscholar.org/CorpusID:247594470>
11. Nguyen, T.H., Nguyen, H.H., Ahmadi, Z., Hoang, T.A., Doan, T.N.: On the impact of dataset size:a twitter classification case study. In: IEEE/WIC/ACM

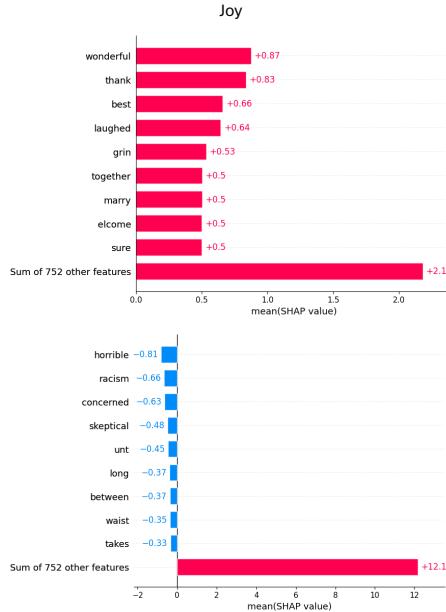
International Conference on Web Intelligence and Intelligent Agent Technology. p. 210–217. WI-IAT '21, Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3486622.3493960>, <https://doi.org/10.1145/3486622.3493960>

12. Qu, Y., Shen, D., Shen, Y., Sajeev, S., Han, J., Chen, W.: Coda: Contrast-enhanced and diversity-promoting data augmentation for natural language understanding (2020), <https://arxiv.org/abs/2010.08670>
13. Rogulsky, S., Popovic, N., Färber, M.: The effects of hallucinations in synthetic training data for relation extraction. In: Joint proceedings of the 2nd workshop on Knowledge Base Construction from Pre-Trained Language Models (KBC-LM 2024) and the 3rd challenge on Language Models for Knowledge Base Construction (LM-KBC 2024) co-located with the 23rd International Semantic Web Conference (ISWC 2024), Baltimore, USA, November 12, 2024 (2024), <https://ceur-ws.org/Vol-3853/paper4.pdf>
14. Stickland, A.C., Murray, I.: Bert and pals: Projected attention layers for efficient adaptation in multi-task learning (2019), <https://arxiv.org/abs/1902.02671>
15. Sun, C., Qiu, X., Xu, Y., Huang, X.: How to fine-tune bert for text classification? (2020), <https://arxiv.org/abs/1905.05583>
16. Susanti, Y.: Data augmentation techniques for process extraction from scientific publications (2024), <https://arxiv.org/abs/2405.14594>
17. Susanti, Y., Färber, M.: Knowledge graph structure as prompt: Improving small language models capabilities for knowledge-based causal discovery. In: The Semantic Web – ISWC 2024: 23rd International Semantic Web Conference, Baltimore, MD, USA, November 11–15, 2024, Proceedings, Part I. p. 87–106. Springer-Verlag, Berlin, Heidelberg (2024). [https://doi.org/10.1007/978-3-031-77844-5\\_5](https://doi.org/10.1007/978-3-031-77844-5_5), [https://doi.org/10.1007/978-3-031-77844-5\\_5](https://doi.org/10.1007/978-3-031-77844-5_5)
18. Verma, P.: Attention is all you need? good embeddings with statistics are enough:large scale audio understanding without transformers/ convolutions/ berts/ mixers/ attention/ rnns or .... (2021), <https://api.semanticscholar.org/CorpusID:238419336>
19. Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Brew, J.: Huggingface's transformers: State-of-the-art natural language processing. CoRR **abs/1910.03771** (2019), <http://arxiv.org/abs/1910.03771>
20. Wolfe, C.R., Lundgaard, K.T.: Exceeding the limits of visual-linguistic multi-task learning. ArXiv **abs/2107.13054** (2021), <https://api.semanticscholar.org/CorpusID:236469329>
21. Xie, Q., Dai, Z., Hovy, E., Luong, M.T., Le, Q.V.: Unsupervised data augmentation for consistency training (2020), <https://arxiv.org/abs/1904.12848>

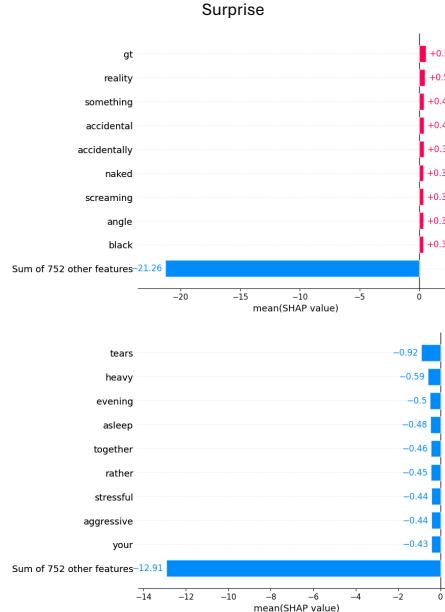
## Appendix

The following shows the class-specific influence graph for each sentiment class produced by SHAP analysis. As explained in §4.2, these visualizations reveal the key features that influence the predictions of the model for each class, as described below.

In the “Joy” class, *wonderful* and *thank* are a strong positive predictor, whereas *horrible* and *racism* are a strong negative predictor. For the “Surprise” class, *gt*<sup>4</sup> followed by *reality*, *something*, and *accidental* appear as a strong predictor, while *tears* and *heavy* have minimal influence. These findings indicate that the model has learned to associate these tokens with class-specific expressions present in the dataset. The influence patterns from the SHAP analysis serve as baseline indicators of the model’s interpretability.



**Fig. 10.** SHAP graph for “Joy” emotion class



**Fig. 11.** SHAP graph for “Surprise” emotion class

<sup>4</sup> Here, *gt* is an artifact from scraped data (representing the “greater than” symbol in HTML encoding), likely introduced during web scraping. Its presence is spurious and should be disregarded in the interpretation.