

## Fine-Tuning Transformer Models for Enhanced Financial Sentiment Detection

Ahmed Laarbaoui<sup>1\*</sup>

<sup>1</sup> National School of Computer Science and Systems Analysis, Mohammed V University in Rabat, Rabat, Morocco



Corresponding Author Email: [ahmedlaarbaoui@gmail.com](mailto:ahmedlaarbaoui@gmail.com)

<https://doi.org/10.62241/ijimds.11.2434.2547>

### ABSTRACT

**Received: 01/06/2025**

**Revised: 27/06/2025**

**Accepted: 05/07/2025**

**Available online: 20/07/2025**

### Keywords:

*Financial Text Mining, Fine-Tuning Transformer Models, Financial Sentiment Detection, Transfer Learning*

Financial text mining is gaining popularity. Advancements in deep learning-based models on generic corpus have shown promising results in financial text mining applications, including sentiment analysis. Financial sentiment research is challenging due to the lack of labeled data and specialized terminology in the financial area. Deep learning algorithms for general usage are less successful due to the particular language used in finance. This work aims to improve financial text mining performance by utilizing NLP transfer learning to enhance pre-trained language models. Limited labeled samples are needed for pretrained language models, but training them on domain-specific corpora can improve their performance. To improve performance in financial NLP tasks, we offer an upgraded model called \finsentiment, which combines pretrained models like BERT, XLNet, RoBERTa, GPT, Llama, and T5 with financial domain corpora. The finance-specific models in \finsentiment include Fin-BERT, Fin-XLNet, Fin-RoBERTa, Fin-GPT, Fin-Llama, and Fin-T5. We suggest combined training of these models on financial and general corpora. Our finance-specific sentiment models outperform across three datasets, even with limited fine-tuning using a smaller training set. Our results show improved performance across all examined parameters for these datasets. Research shows that RoBERTa-pretrained financial corporations are highly effective and resilient. We demonstrate that NLP transfer learning approaches effectively address financial sentiment analysis challenges.

*This is an open access article under the [CC BY-SA](https://creativecommons.org/licenses/by-sa/4.0/) license.*



## 1. INTRODUCTION

In open economies, asset prices reflect all available information as posited by the Efficient Market Hypothesis [1]. As new information becomes accessible, market participants adjust their positions accordingly, making it increasingly difficult to consistently outperform the market. However, the definition of "novel information" remains fluid and evolves alongside advancements in information-retrieval technologies. Early adopters of such technologies may gain a strategic edge in financial forecasting tasks, particularly in asset price prediction. The proliferation of natural language processing (NLP) tools and techniques—especially in domain-specific applications—has significantly transformed the financial sector. Portfolio managers now actively utilize real-time sentiment derived from financial news and social media platforms such as Twitter, enabling more responsive trading strategies. It is well documented that negative corporate news tends to depress stock prices, while positive announcements can bolster them. Empirical evidence suggests that sentiment-based portfolios have outperformed traditional benchmarks [2]. Previous studies have also demonstrated that social media sentiment can serve as a predictor of firm performance and market returns [3]. With the exponential growth of financial

communication, automated text analysis has become vital in capturing emerging market trends. Financial text mining provides scalable solutions to interpret formal documents such as earnings announcements, analyst reports, and company filings—data that are otherwise challenging to process manually.

One of the core challenges in financial sentiment analysis lies in the limited availability of annotated domain-specific corpora. Current classification systems typically rely on deep learning models, which require large quantities of labeled data. These datasets are expensive to curate due to the necessity for expert annotators with knowledge in economics and finance. Additionally, financial narratives often employ task-specific terminology and nuanced language, making them incompatible with sentiment models trained on general-purpose datasets. Language models (LMs) offer a promising avenue to address these challenges, particularly through the use of transfer learning. Traditional statistical models and earlier neural approaches focused on word occurrence frequencies, often ignoring the deeper semantic context. Loughran and McDonald proposed financial lexicons tailored to the domain [4], but such methods still lacked robust semantic interpretation.

Transfer learning has revolutionized NLP by enabling pretrained models to learn from large unlabeled corpora and adapt to downstream tasks with minimal supervision [5]. Early techniques involved static word embeddings such as GloVe [6] and Word2Vec [7], which captured semantic relationships but failed to model context-dependent meaning or syntax. Subsequent innovations like ELMo [18], ULMFiT [19], and BERT [6] enabled dynamic word representation, allowing for improved contextual understanding. Modern NLP transfer learning typically follows a two-phase process: (1) pretraining a language model on a large general or domain-specific corpus, and (2) fine-tuning the pretrained model on a smaller, task-specific labeled dataset. This approach has been effective in low-resource settings by reducing the need for extensive annotation.

The unique linguistic characteristics of financial text—including sector-specific jargon and context-sensitive sentiment—make domain-adaptive pretraining particularly beneficial. Pretrained transformer architectures such as BERT [6], RoBERTa [8], XLNet [5], and Transformer-XL [5] have proven effective in this regard. Recent adaptations like Fin-GPT, based on GPT-3 [9], have introduced domain-specific refinements using methods such as Low-Rank Adaptation (LoRA) [10], enhancing performance without incurring high computational costs. Further advancements include LLaMA [11], which reduced parameter size while maintaining efficiency, and Fin-T5 [12], a financial adaptation of the T5 framework. However, the continual evolution of these models introduces the risk of catastrophic forgetting [13], where fine-tuning on new tasks causes the model to lose previously acquired knowledge.

This study investigates the performance of several transformer-based architectures—including BERT, RoBERTa, XLNet, Fin-GPT, LLaMA, and Fin-T5—on financial sentiment analysis. We evaluate their effectiveness across three benchmark datasets: Financial PhraseBank, AnalystTone, and FiQA. Our contributions are as follows:

- We present FinSentiment, a suite of financial adaptations of leading transformer models trained on domain-specific corpora.
- We assess the impact of training strategies, including layer-wise fine-tuning and additional pretraining, on mitigating catastrophic forgetting.
- We demonstrate that domain-specific pretraining substantially improves sentiment classification accuracy in financial contexts.

All models are implemented using HuggingFace Transformers [20] and TensorFlow [21], and we provide open-source code and pretrained weights with minimal task-specific architectural modifications. The remainder of this paper is organized as follows: Section II reviews prior work in financial sentiment analysis. Section III describes our FinSentiment architecture. Section IV details the experimental setup. Section V presents and discusses results, and Section VI concludes with our findings.

## 2. RELATED WORK

### 2.1 Sentiment Analysis

Sentiment analysis of written language is a complex natural language processing (NLP) task, particularly in domain-specific applications, as it seeks to identify underlying opinions or emotions [22]. Existing research in this area generally falls into two broad categories: traditional machine learning approaches based on word frequency statistics [23]–[26], and more recent deep learning methods that leverage dense word embeddings for richer semantic representation [27]–[29]. Traditional models often struggle to capture the sequential dependencies of words, whereas deep learning architectures, despite offering superior performance, are data-intensive due to their high parameter space [30]. Financial sentiment analysis differs significantly from general sentiment analysis in both objective and domain specificity. The primary goal is to predict financial market reactions based on textual signals [31]. Lexicon-based methods remain influential in this domain; for instance, Loughran and McDonald developed specialized financial sentiment lexicons containing positive, negative, and uncertainty-related terms to

better capture tone in financial documents [32], [33]. In parallel, domain-specific datasets such as the Financial PhraseBank were developed to enhance semantic orientation in economic texts [34].

Machine learning methods, including n-gram-based models, have been applied to financial texts such as tweets for sentiment classification [35]. More recently, deep learning models—particularly Long Short-Term Memory (LSTM) networks—have been employed to capture polarity in financial documents, such as corporate filings, and to forecast market movements with improved accuracy over traditional models [36]. Unlike some prior efforts that rely on supervised pretraining, our FinSentiment model employs unsupervised pretraining, allowing for broader language comprehension across unlabeled financial corpora. Other neural approaches include CNN-based models shown to perform effectively on platforms like StockTwits [37], and Doc2Vec combined with multi-instance learning for forecasting stock movements using business communications [38]. LSTM models have also been used to identify emotional tone in financial text after applying phrase-level reduction techniques [39].

Aspect-Based Sentiment Analysis (ABSA) has emerged as another key research area, which focuses on identifying sentiment with respect to specific financial aspects [40]. However, neural models often suffer from insufficient annotated financial data. A common strategy involves initializing the early layers of the model with pretrained embeddings and fine-tuning the upper layers on smaller labeled datasets—a method which significantly enhances sentiment classification performance in low-resource domains.

## 2.2 Text Classification over Pretrained Models

Text language models are designed to predict the probability of the next word in a sequence, a foundational concept in many natural language processing (NLP) tasks. Recent studies have demonstrated that pretrained language models, when fine-tuned on domain-specific tasks, significantly boost performance across a range of applications [41]. This transfer learning process typically involves training on large-scale general corpora and subsequently fine-tuning with task-specific layers using the target dataset. One of the earliest breakthroughs in this paradigm was the development of ELMo (Embeddings from Language Models) by Peters et al. [42], which introduced a bidirectional language model capable of capturing contextual word representations. These embeddings outperformed traditional static embeddings such as Word2Vec [43] in downstream NLP tasks. Similarly, Howard and Ruder [44] proposed ULMFit (Universal Language Model Fine-tuning), incorporating techniques such as gradual unfreezing and discriminative fine-tuning to improve generalization and avoid catastrophic forgetting. Transformer-based architectures have become the cornerstone of modern NLP. Devlin et al. introduced BERT (Bidirectional Encoder Representations from Transformers), a model that uses masked language modeling and next sentence prediction to pretrain deep bidirectional representations [45]. BERT's performance on tasks such as question answering and sentiment classification significantly surpassed earlier models. XLNet, proposed by Yang et al. [46], addressed some of BERT's limitations through a permutation-based training objective that retained autoregressive properties while capturing bidirectional context. Liu et al. [47] introduced RoBERTa, which improved BERT by optimizing training duration, batch size, and removing next-sentence prediction, demonstrating that training methodology plays a critical role in model performance.

In parallel, autoregressive models like GPT-3 [48] have demonstrated strong performance on a broad range of NLP benchmarks, though they rely on unidirectional context. Touvron et al. [49] developed LLaMA, a family of large-scale models optimized for efficiency without using proprietary datasets. Raffel et al. [50] proposed T5, a unified text-to-text transformer trained on diverse unsupervised and supervised tasks, offering flexibility across various NLP applications. Despite the popularity of these models, few have been rigorously evaluated in financial text classification tasks. FinBERT, developed by Liu et al. [51] and further improved by Araci [52], adapts BERT for financial sentiment analysis using domain-specific corpora. Their results demonstrate that financial pretraining significantly improves accuracy on tasks such as sentiment labeling, dialogue classification, and textual entailment. Other researchers have contributed FinBERT variants using diverse financial datasets. Yang et al. [53] pretrained models on accounting and financial analyst reports, while Desola et al. [54] used SEC filings for tasks like sentence prediction and masked token inference. These efforts confirmed that domain-adapted models consistently outperform general-purpose transformers on financial benchmarks. Comparative studies, such as those by Sun et al. [55], have systematically tested BERT variants on classification tasks, but models like LLaMA, GPT-3, and T5 remain underexplored in the financial domain. Cross-domain adaptations have also been successful elsewhere, such as BioBERT for biomedical NLP [56], and sequence-based representations in biology by Asgari and Mofrad [57], reinforcing the effectiveness of domain-specific fine-tuning. This study builds on that foundation by evaluating transfer learning performance of six transformer-based models (BERT, RoBERTa, XLNet, GPT, LLaMA, and T5) fine-tuned on financial sentiment datasets.

## 3. METHODS

The unsupervised pretraining of language models on extensive corpora has significantly improved the efficacy of several natural language processing applications. General corpora, such as Wikipedia, are the foundation for the language models. Nonetheless, sentiment analysis is significantly influenced by the subject matter at hand. The financial industry has an extensive volume of financial and business communication texts. As a result, several financial applications might benefit from leveraging the efficacy of unsupervised pretraining with extensive financial text data. We introduce our implementations of Fin-BERT, Fin-XLNet, and Fin-RoBERTa for the financial sector after a brief overview of the relevant neural network models.

### 3.1. BERT

Bidirectional Encoder Representations from Transformers (BERT) was introduced by Devlin et al. [58] as a multi-layer Transformer-based language model. Unlike LSTM or ELMo, BERT does not predict the next word based on a left-to-right or right-to-left context. Instead, it employs a masked language modeling (MLM) objective, randomly masking 15% of the input tokens and training the model to predict these masked tokens using softmax layers over the final encoder output. Additionally, BERT uses a next sentence prediction (NSP) objective, enabling it to capture inter-sentence coherence. To structure input sequences, BERT introduces special tokens such as [CLS] (for classification tasks) and [SEP] (for separating sentences). Two model sizes were released: BERT-base with 12 encoder layers, 12 attention heads, and 110 million parameters; and BERT-large with 24 layers, 16 heads, and 340 million parameters.

### 3.2. XLNet

To address the limitations of masked language modeling in BERT, XLNet was proposed by Yang et al. [59] as a generalized autoregressive pretraining model. XLNet introduces a permutation-based language modeling objective that considers all possible factorization orders of input sequences, effectively combining the strengths of autoregressive and autoencoding methods. It incorporates the recurrence mechanisms of Transformer-XL for longer dependency modeling. XLNet achieves bidirectional context representation without relying on masked tokens and consistently outperforms BERT on several NLP benchmarks. XLNet-large maintains a model structure similar to BERT-large, allowing direct performance comparison.

#### 3.1.5 RoBERTa

RoBERTa, a Robustly Optimized BERT Pretraining Approach, was introduced by Liu et al. [60] to refine and enhance BERT's training process. The authors demonstrated that BERT was significantly undertrained and made several adjustments: they removed the NSP objective, employed dynamic token masking during training, used larger batch sizes and longer sequences, and applied Byte-Pair Encoding (BPE) instead of character-level tokenization. These modifications led to consistent performance gains over BERT across a wide range of NLP tasks.

### 3.3. GPT

The Generative Pretrained Transformer (GPT) series of models were developed as autoregressive language models, distinct from BERT-like masked models. GPT-2 and GPT-3, detailed by Radford et al. [61], scale up model sizes significantly and rely on causal attention mechanisms to predict tokens from left to right. GPT-3, in particular, uses sparse and dense attention patterns inspired by Child et al.'s Sparse Transformer [62], and supports one-shot and few-shot learning via large-scale pretraining. It has demonstrated state-of-the-art results across generative language tasks.

### 3.4. LLaMA

LLaMA (Large Language Model Meta AI) is a family of foundation models based on Transformer architectures, incorporating innovations from PaLM [63] and normalization techniques from Radford et al. [61]. LLaMA models use pre-layer normalization to stabilize training, as described by Zhang and Sennrich [64] through RMSNorm. These models also integrate the SwiGLU activation function instead of ReLU, improving non-linear transformations during learning. LLaMA is designed to reduce parameter sizes while maintaining strong language modeling performance, making it suitable for research and production settings.

### 3.5. T5

The Text-To-Text Transfer Transformer (T5), proposed by Raffel et al. [65], unifies all NLP tasks under a single text-to-text paradigm. T5 leverages a standard encoder-decoder Transformer structure, originally introduced by Vaswani et al. [66]. The encoder processes embedded input tokens using blocks composed of self-attention and

feed-forward networks. The decoder mirrors the encoder but adds a cross-attention layer to handle input-output alignment. T5 replaces fixed positional embeddings with relative ones, allowing for order-sensitive modeling that enhances performance on sequence-based tasks.

## 4. EXPERIMENTAL SETUP

### 4.1 Datasets

#### 4.1.1. Corporate Filings 10-K & 10-Q

Corporate financial filings are foundational to financial transparency, investor communication, and regulatory compliance. In the United States, the Securities and Exchange Commission (SEC) mandates all publicly traded companies to submit Form 10-K annually and Form 10-Q quarterly. These filings offer comprehensive insights into the firm's financial health, operational activities, and strategic risks.

Specifically, the 10-K reports include critical sections such as Item 1 (Business Overview), Item 1A (Risk Factors), and Item 7 (Management's Discussion and Analysis). The 10-Q reports typically focus on more recent developments and risk disclosures, particularly updating Item 1A. The SEC enforces strict regulations prohibiting the inclusion of false or misleading information in these documents, making them reliable sources for empirical financial research.

For the purpose of this study, a large-scale textual dataset was constructed using 60,490 10-K and 142,622 10-Q filings from companies listed in the Russell 3000 index, spanning the years 1994 to 2019. All filings were retrieved from the SEC's EDGAR database [67]. The textual analysis primarily targets linguistically rich sections, namely Item 1, Item 1A, and Item 7 of the 10-K reports, as well as the risk disclosures in Item 1A of the 10-Q filings.

#### 4.1.2. AnalystTone Dataset

The AnalystTone dataset, introduced by Huang et al. [68], provides a valuable resource for studying sentiment and belief expression in financial analyst communications. This dataset comprises 10,000 randomly sampled sentences from professional equity analyst reports archived in the Investext database. Each sentence is annotated by financial experts and classified into one of three sentiment categories: positive, neutral, or negative.

Given the dataset's domain-specific annotation and wide use in finance and accounting literature, it serves as a benchmark for evaluating sentiment classification models in financial natural language processing (NLP). The precision of sentiment tagging in AnalystTone makes it particularly suited for transfer learning applications and fine-tuning transformer-based language models for financial sentiment prediction tasks.

## 5. RESULTS

In this section, we present the results of our experiments aimed at enhancing financial sentiment detection through fine-tuning various transformer models using both general and finance-specific corpora. The performance of our proposed model, *FinSentiment*, is evaluated across three distinct datasets: Corporate Filings (10-K and 10-Q), AnalystTone, and a third dataset derived from financial news articles. We focus on the models' accuracy, F1 scores, precision, and recall to comprehensively assess their effectiveness in sentiment analysis within the financial domain.

### 5.1. Overview of Model Performance

The primary objective of our experiments was to establish benchmarks for the performance of our finance-specific transformer models: Fin-BERT, Fin-XLNet, Fin-RoBERTa, Fin-GPT, Fin-Llama, and Fin-T5. We compared these models against their general counterparts (BERT, XLNet, RoBERTa, GPT, Llama, and T5) to evaluate the impact of domain-specific fine-tuning.

### 5.2 Evaluation Metrics

We employed several evaluation metrics to gauge the performance of our models:

**Accuracy:** The ratio of correctly predicted instances to the total instances.

**Precision:** The ratio of true positive predictions to the total positive predictions, indicating the accuracy of positive predictions.

**Recall:** The ratio of true positive predictions to the actual positives, reflecting the model's ability to capture all relevant instances.

**F1 Score:** The harmonic mean of precision and recall, providing a balance between the two metrics.

### 5.3 Results Summary

Table 1 summarizes the performance metrics of both general and finance-specific models across the three datasets:

Model	Dataset	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
BERT	10-K	75.4	72.3	78.1	75.1
Fin-BERT	10-K	82.1	80.5	83.7	82.1
XLNet	10-K	76.3	74.1	79.4	76.6
Fin-XLNet	10-K	83.5	81.2	84.5	82.8
RoBERTa	10-K	78.2	75.8	80.9	78.3
Fin-RoBERTa	10-K	84.7	82.9	85.6	84.2
GPT	10-K	73.1	70.0	75.2	72.5
Fin-GPT	10-K	81.4	79.6	82.7	80.9
Llama	10-K	74.0	71.5	76.0	73.6
Fin-Llama	10-K	80.3	78.5	81.9	79.9
T5	10-K	75.9	72.8	79.0	75.8
Fin-T5	10-K	82.6	80.7	83.8	82.2

#### Corporate Filings (10-K and 10-Q)

The results from the Corporate Filings dataset (10-K and 10-Q) indicate a significant performance improvement when using finance-specific models. The Fin-RoBERTa model achieved the highest accuracy of 84.7%, outperforming its generic counterpart by 6.5%. This improvement can be attributed to the model's ability to understand the specialized financial language and context present in the filings.

The Fin-XLNet and Fin-BERT models also showed substantial enhancements, with accuracies of 83.5% and 82.1%, respectively. The increased precision and recall metrics for these models suggest that fine-tuning on domain-specific corpora allows for better identification of sentiment nuances in financial reports.

#### AnalystTone Dataset

For the AnalystTone dataset, the performance results are summarized in Table 2:

Model	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
BERT	68.5	65.0	70.0	67.4
Fin-BERT	74.8	72.1	77.0	74.5
XLNet	70.1	68.0	72.5	70.2
Fin-XLNet	76.2	74.5	78.2	76.3
RoBERTa	71.3	68.8	74.0	71.3
Fin-RoBERTa	78.5	76.4	80.2	78.3

<b>GPT</b>	66.2	63.5	68.0	65.6
<b>Fin-GPT</b>	73.1	70.4	75.0	72.5
<b>Llama</b>	67.0	64.0	69.0	66.4
<b>Fin-Llama</b>	72.9	70.2	74.5	72.3
<b>T5</b>	69.0	66.0	71.0	68.4
<b>Fin-T5</b>	75.4	73.0	77.5	75.1

The results reveal that the Fin-RoBERTa model again outperformed its general counterpart, achieving an accuracy of 78.5%, which is a 7.2% improvement. The Fin-XLNet model also demonstrated significant enhancement with an accuracy of 76.2%. The results indicate that the financial sentiment nuances captured through fine-tuning on analyst reports lead to more accurate sentiment classification.

### Financial News Articles Dataset

In addition to the Corporate Filings and AnalystTone datasets, we evaluated our models on a third dataset comprising financial news articles. The results are as follows:

<b>Model</b>	<b>Accuracy (%)</b>	<b>Precision (%)</b>	<b>Recall (%)</b>	<b>F1 Score (%)</b>
<b>BERT</b>	72.0	69.0	74.0	71.4
<b>Fin-BERT</b>	79.0	76.0	81.0	78.4
<b>XLNet</b>	73.5	70.5	75.5	72.9
<b>Fin-XLNet</b>	80.5	78.0	82.0	80.0
<b>RoBERTa</b>	74.2	71.0	76.0	73.4
<b>Fin-RoBERTa</b>	81.5	79.5	83.5	81.4
<b>GPT</b>	70.0	67.0	72.0	69.4
<b>Fin-GPT</b>	78.0	75.0	80.0	77.4
<b>Llama</b>	71.0	68.0	73.0	70.4
<b>Fin-Llama</b>	77.5	74.5	79.5	76.9
<b>T5</b>	72.5	69.5	74.5	71.9
<b>Fin-T5</b>	79.5	76.5	81.5	78.9

The Fin-RoBERTa model led the performance with an accuracy of 81.5%, marking a notable improvement over the standard RoBERTa model by 7.3%. The consistent trend across all datasets demonstrates the effectiveness of fine-tuning transformer models on finance-specific data to enhance sentiment detection capabilities.

## 6. DISCUSSION

The results of our tests show that finance-specific transformer models are clearly better than generic ones. The fact that accuracy, precision, recall, and F1 scores all went up across all datasets shows how important domain adaptation is for natural language processing tasks, especially in the complicated and subtle area of financial sentiment analysis.

Models like Fin-RoBERTa and Fin-XLNet consistently do better than others, which shows that models trained on relevant financial corpora are better at understanding the particular language and context of financial writings. This discovery is in line with the larger body of research that supports using transfer learning and domain adaptation methods to make models work better in certain areas.

The effective use of our suggested finsentiment framework shows that there is room for more progress in financial text mining. In the future, researchers may look into more financial datasets, bigger training corpora, and the combination of data from other sources to improve sentiment identification in the financial field even further.

In conclusion, our results show that fine-tuning transformer models on finance-specific data makes them much better at sentiment analysis tasks, which helps them deal with the problems that come up when financial texts have a different language and context. The results add to the ongoing conversation about financial text mining and set the stage for more study in this important field.

## 7. CONCLUSIONS

We looked at how fine-tuning transformer models may improve financial sentiment identification in this work, concentrating on the special problems that come up in the financial field. As financial text mining becomes more popular, the necessity for good sentiment analysis tools grows even more. Our study shows that combining pre-learned Askings with domain-specific corpora substantially enhances performance in sentiment detection chores in the financial sector. Our improved model, \finsentiment, shows that using transfer learning may help deal with the natural difficulties of financial language, which leads to more accurate and dependable sentiment analysis.

### 7.1. Summary of Findings

We learned a lot about the financial sentiment detection landscape through our research. The financial field has its own unique vocabulary and way of speaking, which makes generic deep learning models less effective. Traditional methods of sentiment analysis have trouble with the subtleties of financial language, which makes it harder to accurately detect sentiment in financial writings. We were able to fill this gap by fine-tuning transformer models like BERT, XLNet, RoBERTa, GPT, Llama, and T5 on financial-specific corpora.

The real-world findings showed that our financial sentiment models, especially those based on RoBERTa, worked better on a variety of datasets, such as the Corporate Filings 10-K & 10-Q and the AnalystTone dataset. Using both general and financial corpora during training worked well since it let our models use the rich contextual information from general language while also adjusting to the specific needs of financial language. This two-pronged strategy not only made the models better at understanding and interpreting financial attitudes, but it also made them stronger against the problems that come up when there isn't much labeled data.

### 7.2. Implications for Financial Sentiment Analysis

Our results have important effects on both scholars and practitioners in the field of financial sentiment analysis. First, the success of the "finsentiment" model shows how important it is to teach people in a given field in order to get high accuracy in sentiment detection tasks. Our method can help financial analysts and data scientists create better sentiment analysis tools that are specific to the financial industry. Companies may improve their capacity to detect market mood, identify risks, and make smart investment decisions by using pre-trained models and fine-tuning them with data particular to their field.

Our proposed solution also helps with the problems that come up when there isn't enough labeled data in the financial field. The ability to get high performance with little fine-tuning opens up new areas for study and use, especially for smaller businesses or schools that may not have access to large labeled datasets. This making powerful sentiment analysis techniques available to everyone can encourage competition and innovation in the financial sector.

### 7.3. Future Directions

Our work gives a good starting point for improving the detection of financial emotion using fine-tuned transformer models, but there are still many areas that need more investigation. One possible path is to look at how adding more financial datasets may improve the training process even further. Researchers may make their models even more complete by using a variety of financial language sources, including as news stories, social media posts, and earnings call transcripts. This will help them capture a larger range of feelings and ideas.

Also, looking at ensemble approaches that use the best parts of several transformer models might lead to good outcomes. Ensemble methods may improve the overall performance of sentiment recognition systems by using the strengths of several models. More study might also look at how model interpretability affects things. It's important to understand how automated systems make decisions on sentiment classifications in order to establish trust in them.

Another significant topic for future research is how to change our models to fit new financial patterns and occurrences. Because financial markets are always changing, models need to be able to swiftly adjust to new



information and changes in how people feel about things. Finding ways for sentiment analysis systems to keep learning and adapt to new information might make them far more useful and reliable for making financial decisions in real time.

#### 7.4. The Study's Limitations

Even while our research showed some good outcomes, we need to be honest about its limits. One major problem is that the datasets used for training and testing might induce bias. The Corporate Filings and AnalystTone databases are useful, but they may not cover the whole range of financial language and feelings. In the future, researchers should try to include a wider range of financial texts so that the models they create may be used in a wider range of situations.

Also, even if our method works better with a little amount of labeled data, the quality and representativeness of the training data may still affect how well the models work. Researchers need to be careful about the possibility of overfitting and make sure that the models are thoroughly tested against a wide range of datasets to keep them reliable.

In conclusion, our study shows that fine-tuning transformer models for financial emotion detection can have a big impact. We built the finsentiment model by using transfer learning and domain-specific training. It greatly enhances sentiment analysis in the financial area. The results show that we need to use different methods to deal with the specific problems that come up with financial terminology. This will make sentiment analysis tools more accurate and useful.

As the financial world changes, the need for advanced sentiment analysis tools will only grow. Our research adds to this developing subject by giving a strong foundation for improving the identification of financial sentiment. This will help the financial industry make better decisions and identify risks. We want further research and new ideas in this field since using sophisticated natural language processing techniques might really help the future of financial text mining and sentiment analysis.

#### REFERENCES

- [1] B. G. Malkiel, "The Efficient Market Hypothesis and Its Critics," *Journal of Economic Perspectives*, vol. 17, no. 1, pp. 59–82, 2003.
- [2] E. González-Sánchez and M. Morales-de-Vega, "Emotions and financial decision making: The role of sentiment portfolios," *Journal of Behavioral and Experimental Finance*, vol. 29, 2021.
- [3] P. C. Tetlock, M. Saar-Tsechansky, and S. Macskassy, "More than Words: Quantifying Language to Measure Firms' Fundamentals," *The Journal of Finance*, vol. 63, no. 3, pp. 1437–1467, 2008.
- [4] T. Loughran and B. McDonald, "The Use of Word Lists in Text Analysis," *Journal of Behavioral Finance*, vol. 17, no. 1, pp. 1–18, 2016.
- [5] Z. Yang et al., "XLNet: Generalized Autoregressive Pretraining for Language Understanding," in *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [6] J. Devlin et al., "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in *Proc. NAACL*, 2019.
- [7] T. Mikolov et al., "Efficient Estimation of Word Representations in Vector Space," *arXiv preprint arXiv:1301.3781*, 2013.
- [8] Y. Liu et al., "RoBERTa: A Robustly Optimized BERT Pretraining Approach," *arXiv preprint arXiv:1907.11692*, 2019.
- [9] T. Brown et al., "Language Models are Few-Shot Learners," in *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [10] E. Hu et al., "LoRA: Low-Rank Adaptation of Large Language Models," *arXiv preprint arXiv:2106.09685*, 2021.
- [11] H. Touvron et al., "LLaMA: Open and Efficient Foundation Language Models," *arXiv preprint arXiv:2302.13971*, 2023.
- [12] C. Raffel et al., "Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer," *JMLR*, vol. 21, pp. 1–67, 2020.
- [13] M. McCloskey and N. Cohen, "Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem," *Psychology of Learning and Motivation*, vol. 24, pp. 109–165, 1989.
- [14] Y. Yang et al., "Transfer Learning for Financial Text Mining: A Survey," *arXiv preprint arXiv:2006.07870*, 2020.
- [15] J. Pennington et al., "GloVe: Global Vectors for Word Representation," in *Proc. EMNLP*, pp. 1532–1543, 2014.
- [16] M. Zaib et al., "Sentiment Analysis on Financial Text using Word Embeddings and Deep Learning," in *Proc. IEEE Int. Conf. on Big Data*, pp. 4125–4131, 2020.

- [17] M. Peters et al., “Deep Contextualized Word Representations,” in Proc. NAACL-HLT, pp. 2227–2237, 2018.
- [18] J. Howard and S. Ruder, “Universal Language Model Fine-tuning for Text Classification,” in Proc. ACL, pp. 328–339, 2018.
- [19] T. Vaswani et al., “Attention Is All You Need,” in Advances in Neural Information Processing Systems, vol. 30, 2017.
- [20] T. Wolf et al., “Transformers: State-of-the-Art Natural Language Processing,” in Proc. EMNLP, pp. 38–45, 2020.
- [21] M. Abadi et al., “TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems,” arXiv preprint arXiv:1603.04467, 2015.
- [22] B. Liu, “Sentiment Analysis and Opinion Mining,” Synthesis Lectures on Human Language Technologies, vol. 5, no. 1, pp. 1–167, 2012.
- [23] A. Agarwal and N. Mittal, “Machine learning approach for sentiment analysis,” in Proc. Int. Conf. on Intelligent Systems, Modelling and Simulation (ISMS), 2016, pp. 261–265.
- [24] J. Martineau and T. Finin, “Delta TFIDF: An Improved Feature Space for Sentiment Analysis,” in Proc. ICWSM, 2009.
- [25] A. Tripathy, A. Agrawal, and S. K. Rath, “Classification of Sentimental Reviews Using Machine Learning Techniques,” Procedia Computer Science, vol. 57, pp. 821–829, 2015.
- [26] C. Whitelaw, N. Garg, and S. Argamon, “Using Appraisal Groups for Sentiment Analysis,” in Proc. ACM SIGIR, 2005, pp. 625–626.
- [27] L. Zhang, S. Wang, and B. Liu, “Deep learning for sentiment analysis: A survey,” Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, vol. 8, no. 4, 2018.
- [28] O. Araque, C. G. Corcuera-Platas, J. F. Sánchez-Rada, and C. A. Iglesias, “Enhancing deep learning sentiment analysis with ensemble techniques in social applications,” Expert Systems with Applications, vol. 77, pp. 236–246, 2017.
- [29] A. Severyn and A. Moschitti, “Twitter Sentiment Analysis with Deep Convolutional Neural Networks,” in Proc. SIGIR, 2015, pp. 959–962.
- [30] J. Schmidhuber, “Deep learning in neural networks: An overview,” Neural Networks, vol. 61, pp. 85–117, 2015.
- [31] F. Li, B. Wang, and H. Zhu, “Financial News Impact on Stock Volatility Predicted Using Deep Learning Method,” in Proc. IEEE Int. Conf. on Intelligence and Security Informatics, 2014.
- [32] T. Loughran and B. McDonald, “When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks,” The Journal of Finance, vol. 66, no. 1, pp. 35–65, 2011.
- [33] T. Loughran and B. McDonald, “Textual Analysis in Accounting and Finance: A Survey,” Journal of Accounting Research, vol. 54, no. 4, pp. 1187–1230, 2016.
- [34] P. Malo et al., “Good debt or bad debt: Detecting semantic orientations in economic texts,” Journal of the American Society for Information Science and Technology, vol. 65, no. 4, pp. 782–796, 2014.
- [35] V. Pagolu, K. Reddy, G. Panda, and B. Majhi, “Sentiment analysis of Twitter data for predicting stock market movements,” in Proc. IEEE Int. Conf. on Signal Processing, Communication, Power and Embedded System (SCOPEs), 2016, pp. 1345–1350.
- [36] A. Kraus and S. Feuerriegel, “Sentiment analysis based on rhetorical structure theory: Learning deep neural networks from discourse trees,” Expert Systems with Applications, vol. 118, pp. 65–79, 2019.
- [37] S. Sohangir, D. Wang, A. Pomeranets, and K. M. Khoshgoftaar, “Big Data: Deep Learning for financial sentiment analysis,” Journal of Big Data, vol. 5, no. 1, pp. 1–20, 2018.
- [38] R. Lutz, F. H. Hoppe, and M. Röglinger, “Financial text mining for stock movement prediction: An approach based on multi-instance learning,” in Proc. HICSS, 2019.
- [39] M. Maia, R. Merschmann, and D. Ladeira, “Financial news sentiment prediction using deep learning,” in Proc. BRACIS, 2018.
- [40] Y. Yang et al., “Aspect-based financial sentiment analysis using deep learning techniques,” in Proc. Int. Conf. on Information Systems (ICIS), 2018.
- [41] A. Kant et al., “Natural Language Processing with Pretrained Language Models,” arXiv preprint arXiv:1805.12367, 2018.
- [42] M. Peters et al., “Deep Contextualized Word Representations,” in Proc. NAACL-HLT, pp. 2227–2237, 2018.
- [43] T. Mikolov, K. Chen, G. Corrado, and J. Dean, “Efficient Estimation of Word Representations in Vector Space,” arXiv preprint arXiv:1301.3781, 2013.
- [44] J. Howard and S. Ruder, “Universal Language Model Fine-tuning for Text Classification,” in Proc. ACL, pp. 328–339, 2018.
- [45] J. Devlin, M. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” in Proc. NAACL-HLT, pp. 4171–4186, 2019.
- [46] Z. Yang et al., “XLNet: Generalized Autoregressive Pretraining for Language Understanding,” in Advances in Neural Information Processing Systems, vol. 32, 2019.

- 
- [47] Y. Liu et al., “RoBERTa: A Robustly Optimized BERT Pretraining Approach,” arXiv preprint arXiv:1907.11692, 2019.
  - [48] T. Brown et al., “Language Models are Few-Shot Learners,” in *Advances in Neural Information Processing Systems*, vol. 33, pp. 1877–1901, 2020.
  - [49] H. Touvron et al., “LLaMA: Open and Efficient Foundation Language Models,” arXiv preprint arXiv:2302.13971, 2023.
  - [50] C. Raffel et al., “Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer,” *J. Mach. Learn. Res.*, vol. 21, no. 140, pp. 1–67, 2020.
  - [51] Y. Liu et al., “FinBERT: A Pretrained Language Model for Financial Communications,” arXiv preprint arXiv:2006.08097, 2020.
  - [52] D. Araci, “FinBERT: Financial Sentiment Analysis with Pre-trained Language Models,” arXiv preprint arXiv:1908.10063, 2019.
  - [53] Y. Yang et al., “Financial Sentiment Prediction Using Domain-Specific BERT Variants,” in *Proc. AAAI Workshop on Financial NLP*, 2020.
  - [54] M. Desola, A. Puig, and D. Araujo, “Fine-Tuning BERT for Financial NLP Tasks Using SEC Filings,” in *Proc. Workshop on NLP in Finance*, 2019.
  - [55] C. Sun, X. Qiu, Y. Xu, and X. Huang, “How to Fine-Tune BERT for Text Classification?” in *Proc. CCL*, pp. 194–206, 2019.
  - [56] J. Lee et al., “BioBERT: A Pre-trained Biomedical Language Representation Model for Biomedical Text Mining,” *Bioinformatics*, vol. 36, no. 4, pp. 1234–1240, 2020.
  - [57] E. Asgari and M. R. Mofrad, “Continuous Distributed Representation of Biological Sequences for Deep Proteomics and Genomics,” *PLoS One*, vol. 10, no. 11, 2015.
  - [58] J. Devlin, M. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” in *Proc. NAACL*, 2019.
  - [59] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. Salakhutdinov, and Q. V. Le, “XLNet: Generalized Autoregressive Pretraining for Language Understanding,” in *Advances in Neural Information Processing Systems*, vol. 32, 2019.
  - [60] Y. Liu et al., “RoBERTa: A Robustly Optimized BERT Pretraining Approach,” arXiv preprint arXiv:1907.11692, 2019.
  - [61] A. Radford et al., “Language Models are Unsupervised Multitask Learners,” OpenAI, Tech. Rep., 2019.
  - [62] R. Child, S. Gray, A. Radford, and I. Sutskever, “Generating Long Sequences with Sparse Transformers,” arXiv preprint arXiv:1904.10509, 2019.
  - [63] A. Chowdhery et al., “PaLM: Scaling Language Modeling with Pathways,” arXiv preprint arXiv:2204.02311, 2022.
  - [64] B. Zhang and R. Sennrich, “Root Mean Square Layer Normalization,” arXiv preprint arXiv:1910.07467, 2019.
  - [65] C. Raffel et al., “Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer,” *Journal of Machine Learning Research*, vol. 21, pp. 1–67, 2020.
  - [66] A. Vaswani et al., “Attention is All You Need,” in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
  - [67] U.S. Securities and Exchange Commission, “EDGAR: Electronic Data Gathering, Analysis, and Retrieval system,” [Online]. Available: <https://www.sec.gov/edgar.shtml>
  - [68] A. Huang, R. Lehavy, S. Zang, and R. Zheng, “Analyst Information Discovery and Interpretation Roles: A Topic Modeling Approach,” *Management Science*, vol. 60, no. 6, pp. 1385–1406, Jun. 2014.