

Multi-Task Learning for Sentiment and Topic Classification in Social Media Texts

Rachel Stein^{1*}, Max Keller¹

¹TUM School of Computation, Information and Technology, Munich 80333, Germany

* Corresponding Author: r.stein9@uni-koeln.de

Abstract

Social media platforms have emerged as rich sources of textual data, offering insights into public opinion, consumer preferences, and emerging topics. However, extracting meaningful information from such unstructured and noisy content presents considerable challenges. This study proposes a multi-task learning (MTL) framework to simultaneously perform sentiment classification and topic classification on social media texts. By sharing representations across tasks, the model leverages interrelated patterns and dependencies between sentiment and topical content. Experimental results demonstrate that the MTL approach outperforms single-task baselines in both accuracy and generalization, especially in scenarios with limited labeled data. The integration of attention mechanisms and transformer-based encoders further enhances model interpretability and robustness, offering a scalable solution for real-time social media analytics.

Keywords

Multi-Task Learning, Sentiment Classification, Topic Classification, Social Media Text, Transformer, Attention Mechanism, Deep Learning.

1. Introduction

Social media has become a dominant mode of digital communication, with billions of users actively engaging in sharing opinions, experiences, and information across platforms such as Twitter, Reddit, Facebook, and TikTok[1]. These platforms generate massive volumes of textual data in real time, forming a dynamic, continuously evolving stream of public discourse[2]. From political campaigns and brand perception to global crises and entertainment, social media content reflects both individual sentiment and collective awareness[3]. This makes it an invaluable resource for applications in marketing, social research, public health monitoring, and political forecasting[4].

However, analyzing this content presents significant challenges. Social media texts are inherently noisy, characterized by informal language, abbreviations, emojis, misspellings, and short sentence structures[5]. Moreover, user intent and sentiment are often subtle or context-dependent, making conventional natural language processing (NLP) methods insufficiently precise[6]. Extracting actionable insights from such data requires sophisticated models capable of capturing both the emotional tone (sentiment) and the semantic focus (topic) of each message[7].

Sentiment classification seeks to determine the emotional polarity of a piece of text—commonly categorized as positive, negative, or neutral[8]. Topic classification, on the other hand, involves assigning a given text to predefined thematic categories such as politics, sports, health, or technology. Traditionally, these two tasks have been addressed separately using distinct datasets and model pipelines. Yet in practical scenarios, sentiment and topic often influence each other[9]. For example, posts discussing climate change may carry negative

sentiment due to concern or frustration, while those about sports events may skew more positive. Ignoring this interdependence can limit the depth and nuance of automated text understanding[10].

Multi-task learning (MTL) offers a compelling solution to this problem by enabling a model to learn multiple tasks simultaneously while sharing internal representations[11]. The core idea is that by training on related objectives—in this case, sentiment and topic classification—the model can leverage commonalities between tasks, leading to improved generalization and robustness[12]. MTL is particularly advantageous in domains where labeled data is sparse or imbalanced, as learning signals from one task can enhance performance on another[13].

Recent advancements in deep learning, especially the introduction of transformer-based architectures such as Bidirectional Encoder Representations from Transformers (BERT), have revolutionized NLP[14]. These models excel at capturing contextual relationships between words and phrases, making them ideal backbones for MTL frameworks. By integrating attention mechanisms and hierarchical representations, transformer models can disentangle overlapping signals from social media texts, identifying both emotional tone and semantic focus with higher accuracy.

In this study, we propose a unified deep learning architecture that simultaneously performs sentiment and topic classification on social media texts. The model employs a shared transformer encoder to extract context-aware embeddings, followed by task-specific attention layers and output classifiers. This design allows for both global and task-specific feature learning. We evaluate our framework on benchmark social media datasets and conduct detailed ablation studies to quantify the benefits of MTL. Our findings reveal that the proposed approach consistently outperforms single-task models, particularly in low-resource settings and when confronted with ambiguous or multi-topic texts.

By advancing the methodological toolkit for analyzing social media content, this research contributes to the growing field of AI-driven social understanding. The ability to accurately and efficiently interpret public sentiment and thematic concerns in real time has profound implications for policy-making, crisis response, and consumer engagement strategies.

2. Literature Review

The rapid growth of user-generated content on social media platforms has accelerated research in automated text analysis, particularly in the areas of sentiment classification and topic classification[15]. These tasks have traditionally been addressed through separate pipelines, with sentiment analysis focusing on emotional tone detection and topic classification targeting semantic categorization[16]. Early approaches to sentiment classification relied on lexicon-based methods, where predefined word lists with assigned sentiment scores were used to infer polarity[17]. While straightforward, such approaches suffered from poor generalization and limited handling of contextual sentiment shifts[18].

Machine learning-based models, particularly support vector machines and naive Bayes classifiers, marked the first wave of performance improvements in sentiment classification[19]. These models operated on bag-of-words or TF-IDF representations and offered increased robustness over lexicon-based methods[20]. Topic classification followed a similar trajectory, initially adopting rule-based or keyword-matching systems before transitioning to supervised learning frameworks. However, both approaches shared limitations in capturing syntactic and semantic nuance, particularly in the highly informal and context-dependent environment of social media[21].

The introduction of deep learning brought about a paradigm shift in NLP. Recurrent neural networks, and later convolutional architectures, allowed for modeling of word sequences and local phrase patterns, enabling better contextual understanding[22]. In sentiment

classification, deep models demonstrated the ability to identify complex sentiment cues such as sarcasm, irony, or negation. Topic classification similarly benefited from hierarchical models that considered document-level structures and word co-occurrence patterns[23]. Nevertheless, these models still required substantial labeled data and did not inherently model the interdependence between sentiment and topic.

The emergence of transformer-based models has further transformed the NLP landscape. These models, typified by BERT, leverage self-attention mechanisms to encode contextual dependencies across entire text sequences[24]. Their pretraining on massive text corpora enables strong performance on a range of downstream tasks with minimal fine-tuning. In both sentiment and topic classification, transformers have established state-of-the-art benchmarks, particularly when fine-tuned on task-specific datasets[25].

Despite these advancements, most prior work has treated sentiment and topic classification as independent tasks. This separation disregards the natural coupling that exists between emotion and content in real-world communication[26]. Posts on controversial topics may naturally carry strong sentiment, while emotionally neutral expressions often pertain to less subjective themes. Ignoring such associations can lead to misclassifications, especially in ambiguous or short texts typical of social media[27].

MTL addresses this gap by allowing models to jointly learn multiple objectives. The shared representation layer in MTL models captures common linguistic patterns useful for both tasks, while task-specific layers enable fine-grained learning[28]. Research in MTL has shown that auxiliary tasks can provide regularization effects, reducing overfitting and improving generalization. For sentiment and topic classification, this means that emotional polarity can be informed by topic context and vice versa. In low-resource settings, MTL also facilitates knowledge transfer, where high-resource tasks help improve performance on data-scarce counterparts[29].

Several MTL architectures have been proposed, ranging from hard parameter sharing—where early layers are fully shared between tasks—to soft parameter sharing, where task-specific networks are trained with shared regularization constraints[30]. In recent studies involving social media text, models have adopted hybrid attention mechanisms to selectively share features across tasks. This enables the model to focus on relevant linguistic cues for each objective while leveraging common patterns like negation, subjectivity markers, or thematic keywords.

While promising, challenges in MTL remain. Balancing task-specific loss functions, preventing negative transfer, and handling label imbalances are ongoing areas of research. Moreover, the diversity of social media language—ranging from formal discourse to meme-driven slang—requires models that are not only robust but also adaptable to evolving communication styles.

This review highlights the need for a unified framework that simultaneously models sentiment and topic in a way that reflects their natural co-occurrence and mutual influence. By integrating transformer-based representations with MTL strategies, it is possible to create models that deliver deeper and more nuanced understanding of social media texts, enabling a broad range of analytical and predictive applications.

3. Methodology

This section presents the methodology adopted to develop a multi-task learning framework for sentiment and topic classification in social media texts. The process includes data collection and preprocessing, model design, and training configuration.

3.1. Data Collection and Preprocessing

To support the dual-task objective, we compiled a dataset from publicly available Twitter and Reddit posts. Each entry was annotated with a sentiment label (positive, negative, or neutral) and a topic label (e.g., politics, health, entertainment, technology). Raw text data was preprocessed to reduce noise and standardize input. Steps included converting text to lowercase, removing URLs, mentions, hashtags, emojis, and punctuations. We applied a transformer-compatible tokenizer to convert text into tokens, truncating or padding them to a fixed length of 128 tokens to ensure uniformity in input size as in Figure 1.

Additionally, we balanced the dataset to avoid bias toward dominant classes. One-hot encoding was used to represent both sentiment and topic labels for supervised learning. The dataset was split into training, validation, and testing subsets using a stratified sampling approach, maintaining proportional representation across both tasks.




Figure 1. Data Preprocessing Workflow

3.2. Model Architecture

The model architecture centers on a shared encoder and two task-specific output heads. A pre-trained BERT base transformer serves as the shared encoder, capturing the semantic and syntactic patterns within input text. The shared representation generated by BERT is passed simultaneously to two separate branches: one for sentiment classification and the other for topic classification.

Each output head includes a dropout layer to reduce overfitting, followed by a dense layer with a softmax activation function that outputs probabilities for each class as in Figure 2. To better handle task interference, we integrate a lightweight attention-based mechanism that selectively emphasizes different parts of the encoder's output for each task. This enables the model to retain shared knowledge while respecting the unique demands of each classification objective.




Figure 2. Model Architecture

3.3. Training Configuration

The training process aims to jointly optimize both classification tasks using a shared representation. We use the AdamW optimizer with a learning rate of $3e-5$ and apply early stopping based on the macro-average F1 score on the validation set. The model is trained in batches of 32 samples for up to 10 epochs. During training, loss gradients from both tasks are computed and used to update the shared encoder weights, allowing the model to learn generalized features that benefit both outputs.

A warmup schedule is employed for the learning rate, gradually increasing it during the initial steps to stabilize training. Dropout regularization and L2 weight decay are also applied to prevent overfitting. We conduct periodic evaluations every 500 steps using the validation set, tracking precision, recall, and F1 scores for both sentiment and topic classification as in Figure 3.




Figure 3. Training Loop Overview

4. Results and Discussion

This section presents the empirical results of our proposed multi-task learning model and compares its performance against single-task baselines. We evaluate each model on sentiment classification and topic classification using F1 score, a metric that balances precision and recall and is well-suited for imbalanced classification tasks. All experiments were conducted using the same training-validation-test splits to ensure fairness and reproducibility.

4.1. Performance Comparison with Single-Task Baselines

The first baseline model is trained solely on sentiment labels, while the second is optimized only for topic classification. In contrast, our multi-task model jointly learns both tasks using a shared encoder with task-specific heads. The evaluation results demonstrate that multi-task learning offers consistent improvements in both classification objectives.




Figure 4. Performance Comparison Between Models

As in Figure 4, the multi-task model achieves an F1 score of 0.83 on sentiment classification, outperforming the sentiment-only model (0.78). On topic classification, the multi-task model also shows a noticeable improvement, reaching an F1 score of 0.80 compared to the topic-only baseline of 0.74. These results suggest that learning representations jointly for both tasks leads to better generalization and reduces overfitting, especially for underrepresented classes.

4.2. Cross-Task Learning Benefits

Further analysis reveals that the sentiment classifier benefits from the semantic structure introduced by topic learning. For instance, posts associated with political discourse tend to skew negative, while those about entertainment or sports are often positive. The joint model implicitly learns these correlations, improving its ability to resolve ambiguous cases where sentiment is dependent on topic context.

Topic classification also gains from sentiment information, as emotional cues can provide additional context for identifying the correct thematic category. For example, a post expressing fear or anxiety may be more likely to fall under health or crisis-related topics, while excitement may signal content related to entertainment or events.

4.3. Generalization to Unseen Data

The multi-task model demonstrates improved robustness when evaluated on a held-out dataset containing previously unseen topics and informal text structures. Unlike the single-task models, which show a notable performance drop under domain shift, the multi-task model maintains relatively stable F1 scores. This suggests that shared learning helps capture generalized language patterns beyond task-specific cues.

5. Conclusion

In this study, we proposed a multi-task learning framework that jointly performs sentiment and topic classification on social media texts. Social media content is inherently multifaceted, often reflecting both user opinions and thematic concerns simultaneously. Traditional single-task models, while effective to some extent, fall short when it comes to capturing the shared linguistic and semantic patterns that underpin both sentiment and topic information.

By training a shared representation layer that feeds into task-specific classifiers, our multi-task model effectively leverages inter-task relatedness, leading to better generalization and improved performance. Empirical results across multiple evaluation metrics demonstrated that the proposed model consistently outperformed its single-task counterparts in both sentiment and topic classification. The performance gains were especially prominent on noisy, user-generated content, highlighting the robustness of the joint learning approach.

Moreover, our analysis showed that multi-task learning not only enhances classification accuracy but also contributes to more coherent and context-aware predictions. This is particularly valuable in practical applications such as brand monitoring, political discourse analysis, or customer service automation, where understanding both the sentiment and the underlying topic is crucial.

Despite the encouraging results, some limitations remain. Our current framework assumes equal importance of both tasks, which may not always align with specific application needs. Future work could explore adaptive weighting mechanisms to balance task contributions dynamically. Additionally, incorporating external knowledge sources such as domain ontologies or sentiment lexicons could further enhance the interpretability and robustness of the model.

In conclusion, our findings underscore the efficacy of multi-task learning as a strategic solution for complex NLP tasks in social media analytics. The integration of sentiment and topic

modeling within a unified framework offers a powerful tool for more nuanced and insightful text classification, paving the way for broader adoption in real-world applications.

References

- [1] Ovhal, D. R. (2025). EVOLUTION OF SOCIAL MEDIA: AS INNOVATIVE RESEARCH AND EDUCATIONAL TOOL. *The Infinite Path: Research Across Disciplines*, 136.
- [2] Han, X., Yang, Y., Chen, J., Wang, M., & Zhou, M. (2025). Symmetry-Aware Credit Risk Modeling: A Deep Learning Framework Exploiting Financial Data Balance and Invariance. *Symmetry* (20738994), 17(3).
- [3] Mirbabaie, M., Bunker, D., Stieglitz, S., Marx, J., & Ehnis, C. (2020). Social media in times of crisis: Learning from Hurricane Harvey for the coronavirus disease 2019 pandemic response. *Journal of Information Technology*, 35(3), 195-213.
- [4] Igwama, G. T., Olaboye, J. A., Maha, C. C., Ajegbile, M. D., & Abdul, S. (2024). Big data analytics for epidemic forecasting: Policy Frameworks and technical approaches. *International Journal of Applied Research in Social Sciences*, 6(7), 1449-1460.
- [5] Khan, J., Ahmad, K., Jagatheesaperumal, S. K., & Sohn, K. A. (2025). Textual variations in social media text processing applications: challenges, solutions, and trends. *Artificial Intelligence Review*, 58(3), 89.
- [6] Yaghoobian, H. (2021). Structure-and Context-aware NLP Approaches to Emotion and Subjectivity (Doctoral dissertation, University of Georgia).
- [7] George, A. S., & Baskar, T. (2024). Leveraging big data and sentiment analysis for actionable insights: A review of data mining approaches for social media. *Partners Universal International Innovation Journal*, 2(4), 39-59.
- [8] Mohammad, S. M. (2021). Sentiment analysis: Automatically detecting valence, emotions, and other affectual states from text. In *Emotion measurement* (pp. 323-379). Woodhead Publishing.
- [9] Tan, K. L., Lee, C. P., & Lim, K. M. (2023). A survey of sentiment analysis: Approaches, datasets, and future research. *Applied Sciences*, 13(7), 4550.
- [10] Srivastava, S., Chakraborty, C., & Sarkar, M. K. (2024). Leveraging machine learning and dimensionality reduction for sports and exercise sentiment analysis. *Measurement: Sensors*, 33, 101182.
- [11] Zhang, Q., Chen, S., & Liu, W. (2025). Balanced Knowledge Transfer in MTTL-ClinicalBERT: A Symmetrical Multi-Task Learning Framework for Clinical Text Classification. *Symmetry*, 17(6), 823.
- [12] Tarekegn, A. N., Giacobini, M., & Michalak, K. (2021). A review of methods for imbalanced multi-label classification. *Pattern Recognition*, 118, 107965.
- [13] Singla, S., Thakur, A., Swami, A., Sawarn, U., & Singla, P. (2024, November). Advancements in Natural Language Processing: BERT and Transformer-Based Models for Text Understanding. In *2024 Second International Conference on Advanced Computing & Communication Technologies (ICACCTech)* (pp. 372-379). IEEE.
- [14] Liu, Y., Ren, S., Wang, X., & Zhou, M. (2024). Temporal Logical Attention Network for Log-Based Anomaly Detection in Distributed Systems. *Sensors*, 24(24), 7949.
- [15] Rogers, D., Preece, A., Innes, M., & Spasić, I. (2021). Real-time text classification of user-generated content on social media: Systematic review. *IEEE Transactions on Computational Social Systems*, 9(4), 1154-1166.
- [16] Nandwani, P., & Verma, R. (2021). A review on sentiment analysis and emotion detection from text. *Social network analysis and mining*, 11(1), 81.
- [17] Bonta, V., Kumares, N., & Janardhan, N. (2019). A comprehensive study on lexicon based approaches for sentiment analysis. *Asian Journal of Computer Science and Technology*, 8(S2), 1-6.
- [18] Li, P., Ren, S., Zhang, Q., Wang, X., & Liu, Y. (2024). Think4SCND: Reinforcement Learning with Thinking Model for Dynamic Supply Chain Network Design. *IEEE Access*.

- [19] Palanivinayagam, A., El-Bayeh, C. Z., & Damaševičius, R. (2023). Twenty years of machine-learning-based text classification: A systematic review. *Algorithms*, 16(5), 236.
- [20] Tan, Y., Wu, B., Cao, J., & Jiang, B. (2025). LLaMA-UTP: Knowledge-Guided Expert Mixture for Analyzing Uncertain Tax Positions. *IEEE Access*.
- [21] Alshattnawi, S., Shatnawi, A., AlSobeh, A. M., & Magableh, A. A. (2024). Beyond word-based model embeddings: Contextualized representations for enhanced social media spam detection. *Applied Sciences*, 14(6), 2254.
- [22] Wang, J., Tan, Y., Jiang, B., Wu, B., & Liu, W. (2025). Dynamic Marketing Uplift Modeling: A Symmetry-Preserving Framework Integrating Causal Forests with Deep Reinforcement Learning for Personalized Intervention Strategies. *Symmetry*, 17(4), 610.
- [23] Fitz, S., & Romero, P. (2021). Neural networks and deep learning: A paradigm shift in information processing, machine learning, and artificial intelligence. *The Palgrave handbook of technological finance*, 589-654.
- [24] Poria, S., Hazarika, D., Majumder, N., & Mihalcea, R. (2020). Beneath the tip of the iceberg: Current challenges and new directions in sentiment analysis research. *IEEE transactions on affective computing*, 14(1), 108-132.
- [25] Xing, S., Wang, Y., & Liu, W. (2025). Multi-Dimensional Anomaly Detection and Fault Localization in Microservice Architectures: A Dual-Channel Deep Learning Approach with Causal Inference for Intelligent Sensing. *Sensors*.
- [26] Nieminen, M. (2023). The Transformer Model and Its Impact on the Field of Natural Language Processing.
- [27] Gautam, H., Gaur, A., & Yadav, D. K. (2025). A Survey on the Impact of Pre-Trained Language Models in Sentiment Classification Task. *International Journal of Data Science and Analytics*, 1-39.
- [28] Wu, B., Qiu, S., & Liu, W. (2025). Addressing Sensor Data Heterogeneity and Sample Imbalance: A Transformer-Based Approach for Battery Degradation Prediction in Electric Vehicles. *Sensors*, 25(11), 3564.
- [29] Wang, J., Zhang, H., Wu, B., & Liu, W. (2025). Symmetry-Guided Electric Vehicles Energy Consumption Optimization Based on Driver Behavior and Environmental Factors: A Reinforcement Learning Approach. *Symmetry*.
- [30] Ren, S., Jin, J., Niu, G., & Liu, Y. (2025). ARCS: Adaptive Reinforcement Learning Framework for Automated Cybersecurity Incident Response Strategy Optimization. *Applied Sciences*, 15(2), 951.
- [31] Li, P., Ren, S., Zhang, Q., Wang, X., & Liu, Y. (2024). Think4SCND: Reinforcement Learning with Thinking Model for Dynamic Supply Chain Network Design. *IEEE Access*.