

Real-Time Sentiment and Emotion Analysis for Social Media Brand Monitoring Using Transformer-Based Natural Language Processing

Monali Parikh¹ Vaibhavi Pandya² Priyanka Makwana³ Barot Kajal⁴

¹ Assistant Professor, Department of Artificial Intelligence and Machine Learning, Parul Institute of Engineering and Technology, Parul University Vadodara Gujarat

Mail Id: monaliparikh89@gmail.com

² Assistant Professor, Department of Computer Science and Engineering, Krishna school of Diploma Studies, Drs. Kiran and Pallavi Patel Global University, Vadodara, Gujarat.

Mail id: vaibhavipandya95@gmail.com

³ Assistant Professor, Department of Computer Science and Engineering, Krishna school of Engineering and Technology, Drs. Kiran and Pallavi Patel Global University, Vadodara, Gujarat.

Mail id: priyankamakwana2009@gmail.com

⁴ Assistant Professor, Department of Computer Science and Engineering, Krishna school of Engineering and Technology, Drs. Kiran and Pallavi Patel Global University, Vadodara, Gujarat.

Mail id: barotkajal96@gmail.com

Abstract With the explosion of social media, brand perception is no longer a one-way street but can be seen and felt by anyone, which presents an opportunity and an operational risk for organizations. Consumer opinion, intensity of feeling, and rising reputational crises are represented in high volume and high velocity text streams on platforms like X, Instagram, Facebook, Reddit and LinkedIn. Most lexicon-based sentiment models are inadequate for sarcasm, informal speech and variation of context, all of which exist in UGC. This paper introduces an online sentiment and emotion classification model for social media brand monitoring that relies on transformer models for NLP, along with an interactive dashboard for analytics. The architecture consists of four stages: data acquisition, text normalization, contextual representation learning and classification, and visualization on the dashboard. A fine-tuned BERT model is used as the main contextual classification model and compared with baselines of TF-IDF sequence classifiers (Naivebayes, SVM) and a BiLSTM sequence model. In addition to polarity detection, the proposed framework also recognizes emotions based on Ekman categories, which is a multi-class emotion recognition. The criteria for evaluation are accuracy, precision, recall, Macro-F1, inference latency and dashboard responsiveness. Keywords: Sentiment Analysis, Emotion Recognition, Social Media Monitoring, Brand Intelligence, BERT, Transformer Models, Real-time NLP, Affective Computing, Dashboard Analytics.

1. INTRODUCTION

Social media is now playing a key role in brand intelligence, with consumers talking about products and campaigns almost instantly, in ways that are unstructured and can go viral – making monitoring a strategic imperative and not a convenience [1][2]. Nevertheless, the informal language (sarcasm, code-mixing, emojis and negation) makes lexicon-based and shallow ML pipelines unreliable [3][4] and prompts the development of transformer-based models like BERT which can capture deep context-based meaning, but are still limited in terms of interpretability and multilingualism robustness [5][6]. This paper proposes a modular and synthesized real-time architecture for brand monitoring that adopts a reproducible evaluation protocol, thus filling the gap between brand monitoring research and real-time systems for production use.

2. RELATED WORK

The evolution of sentiment analysis includes lexicon-based scoring, supervised feature-based learning and, most recently, deep neural and transformer-based modeling. Yet, hand-crafted dictionaries and sparse representations like bag-of-words or TF-IDF are not always effective for fine-grained classification and have some limitations, like domain shift, negation, or lexical ambiguity, particularly in noisy, fast-changing social media text. Xu et al. point to the difficulties in high content velocity, temporal drift, and sarcasm. The bidirectional pre-training of BERT allows for better performance with short and ambiguous posts, while RoBERTa increases the rigor of the training and XLM-R expands language coverage. A parallel research stream focuses on fine-grained emotion recognition, as the three-class schemes cannot differentiate anger, fear, disappointment and joy. [7].

3. RESEARCH GAP AND CONTRIBUTIONS

The current gap between top performing offline text classifiers and deployable brand monitoring systems is a large one: while many studies report benchmark performance, few consider how to integrate the pipeline, how long it takes, or how interpretable it is. The second is about emotional granularity, and outrages, anxiety and disappointment call for different responses, but frequently they all fall under a negative umbrella. The third one is the lack of a comparative study of latency/accuracy between transformer and classical deployments. This paper proposes a solution to these drawbacks by providing an end-to-end pipeline from acquisition to preprocessing, classification, and visualization in

a dashboard, using dual-layer prediction with sentiment polarity and emotion recognition, and comparing the modeling of accuracy with deployment efficiency in terms of Macro-F1, ROC-AUC, and latency, and discussing deployment issues related to dashboard responsiveness, ingestion throughput, privacy, and model maintenance.

4. PROPOSED METHODOLOGY

The pipeline consists of four modules: data acquisition, text preprocessing, sentiment/emotion classification, and visualization in dashboard.

A. System Architecture

Social media posts are retrieved through platform APIs using entity-centric keyword, hashtag, and brand-mention filters. A message-queue layer (Apache Kafka / RabbitMQ) buffers high-velocity streams to absorb volume bursts during product launches or crisis episodes. The preprocessing and classification layer applies adaptive normalization followed by dual-task BERT inference; a hybrid cascade router forwards high-confidence posts to a lightweight SVM fast path and ambiguous posts to the multi-task BERT deep path, producing sentiment polarity (positive, negative, neutral) and six-class Ekman-style emotion labels (joy, sadness, anger, fear, surprise, disgust). Enriched prediction events are published to MongoDB/InfluxDB and rendered as sentiment panels, Net Sentiment Score trends, emotion distributions, alert notifications, and keyword heat maps, as shown in Fig. 1




Figure 1: Proposed System Architecture — Real-Time Social Media Sentiment & Emotion Monitoring [8]

B. Text Preprocessing

This normalization pipeline excludes URLs and anonymizes any users referenced in posts to suppress non-semantic noise and to maintain the structure of posts. The model-family branch is then subjected to aggressive normalization (removing the stop words, lemmatizing) for classical models based on TF-IDF, and conservative normalization (keeping punctuation, casing and syntactic markers important for the contextual encoding) for BERT based models. Both branches share a common feature for converting emoji to text, decomposing hashtags (e.g. #ProductLaunch → product launch), normalizing spaces (e.g. " hello " → "hello"), and marking the scope of a negation so that it doesn't get inverted (e.g. "I am not a Frog" → "I am not a frog"). The pipeline is implemented as a stateless microservice that allows to scale and hot-swap preprocessing profiles independently [9,10].

C. Feature Extraction

Three families of embeddings are used: (i) TF-IDF vectors for classical baselines; (ii) static embeddings (Word2Vec [11] or GloVe [12] in which transformer inference is computationally limited; and (iii) BERT contextual embeddings as the main source of embeddings, which include subtle emotion cues and sentiment reversal.

5. MODEL LAYER

A common BERT encoder and task-specific output heads are used to jointly implement three-class sentiment classification and six-class emotion recognition, which makes the implementation more efficient and allows for better transfer between related tasks. As depicted in Fig. 2, the input post is first tokenized using WordPiece approach by prepending [CLS] token then encoded by 12 stacked bert-base-uncased transformer layers, generating 768-dimensional contextual representation. The final-layer [CLS] representation h_i is input to a Sentiment Head ($W_s h_i + b_s$, 3-way softmax) and an Emotion Head ($W_e h_i + b_e$, 6-way softmax). The two heads are optimized together with a weighted multi-task cross-entropy loss, and the shared encoder learns representations that are useful for both of the affective tasks.




Fig. 2. Dual-task BERT architecture for joint sentiment and emotion classification.

TABLE I. COMPARATIVE MODEL SUITE

ID	Description	Use
B1	Multinomial NB + TF-IDF	Lexical baseline
B2	Linear SVM + TF-IDF	Classical benchmark
B3	BiLSTM + GloVe	Sequential baseline
P1	Fine-tuned BERT	Contextual transformer
P2	Multi-task BERT	Joint learning model
H1	Hybrid Cascade	Deployment efficiency

6. MATHEMATICAL FORMULATION

Let the input stream of social media posts be $X = \{x_1, x_2, \dots, x_n\}$, where each x_i is a timestamped textual message. For each post x_i , the system predicts a sentiment label $y_i^s \in \{\text{positive, negative, neutral}\}$ and an emotion label $y_i^e \in \{\text{joy, sadness, anger, fear, surprise, disgust}\}$.

The sentiment prediction is: $\hat{y}_i^s = \text{argmax softmax}(W_s h_i + b_s)$, where $h_i = f\theta(x_i)$. The emotion prediction is: $\hat{y}_i^e = \text{argmax softmax}(W_e h_i + b_e)$. For joint multi-task learning: $L = \lambda_s L_s + \lambda_e L_e$, where λ_s and λ_e are task-weighting coefficients. The Net Sentiment Score at time t is: $\text{NSS}(t) = [N_{\text{pos}}(t) - N_{\text{neg}}(t)] / N_{\text{total}}(t)$.

Evaluation metrics: Precision $P = TP / (TP + FP)$; Recall $R = TP / (TP + FN)$; $F1 = 2 \times (P \times R) / (P + R)$. Macro-F1 is the unweighted mean of per-class F1 scores.

7. DATASET STRATEGY

The TweetEval benchmark provides fixed, standardized splits for tweet classification. Its sentiment subset (three labels) contains 45,615 training, 2,000 validation, and 12,284 test instances; its emotion subset (anger, joy, optimism, sadness) contains 3,257 training, 374 validation, and 1,421 test instances. Extension to the full six-category Ekman taxonomy is left as future work. The benchmark for external validation, temporal tracking, and latency testing is complemented by a domain specific validation corpus of brand names, product mentions, and campaign hashtags gathered from X, Reddit and Instagram, following platform terms of service.

TABLE II. TWEETVAL DATASET STATISTICS

Subset	Train	Val.	Test
Sentiment (3-class)	45,615	2,000	12,284
Emotion (4-class)	3,257	374	1,421

8. EXPERIMENTAL SET UP AND EVALUATION

The experiments address four research questions: (1) Can contextual transformers beat traditional baselines for sentiment polarity? (2) Is it possible to get better interpretability using fine-grained emotion classification than just polarity? (3) What are the compromise between predictive performance and cost of operation? (4) Can models trained on a benchmark set of data be used on a domain-specific brand-monitoring corpus? The models are trained using a strict train-validation-test split as in TweetEval [15] and hyperparameter optimization is limited to the validation split. The transformer models use early stopping while linear learning-rate warmup with decay and class-weighted loss are applied and each experiment is conducted with random seeds, reporting the mean and the standard deviation. Evaluation includes accuracy, precision, recall, macro/weighted/per-class F1, ROC-AUC/PR-AUC for alerting tasks, confusion analysis and deployment metrics (per-post inference latency, end-to-end pipeline latency, end-to-end streaming throughput, alert-detection delay).

TABLE III ANTICIPATED PERFORMANCE RANGES

Model	Acc. (%)	Sent. F1	Latency
MNB+TF-IDF	61-65	60-63	<1 ms
SVM+TF-IDF	64-68	62-66	<1 ms
BiLSTM+GloVe	67-71	66-70	5-15 ms
Fine-tuned BERT	73-78	72-77	30-80 ms
Multi-task BERT	74-79	73-78	35-85 ms
Hybrid Cascade	72-77	71-76	5-20 ms

9. CONCLUSION AND FUTURE WORK:

This framework can be used for brand-reputation management, crisis detection, monitoring campaigns, and analysing customer interactions. Moreover, it presents actionable intelligence by considering emotion and sentiment together, employing a cost-aware hybrid cascade. Limitations encompass limited English-only text, problems with sarcasm and mixed feelings, coordination signals, and generic benchmarks in dashboards that don't match the specific needs of a brand, dashboards should display model confidence for human auditing. Some of the future research works are multilingual scaling, multimodal sentiment analysis, topic- and event-aware modeling, online vocabulary adaptation, explainability (LIME/SHAP), and zero-shot monitoring with large language models. This research proposes a real-time transformer-based sentiment and emotion analysis system for the real-time brand monitoring. For effective monitoring, it is important to have an end-to-end architecture including streaming ingestion, text normalisation, contextual and emotion-aware classification and dashboards to help with decision-making. Adding emotion features as a first-class citizen along with sentiment, plus a classic accuracy measure, Macro-F1, ROC-AUC, latency and dashboard responsiveness, brings the gap to deployable brand intelligence. Empirical benchmark validation, multilingual extension and hybrid deployment strategies are planned for the future.

10. REFERENCES

- [1] Twitter Developer Platform, "Twitter API v2: Academic Research Access," Twitter, Inc., 2023.
- [2] Meta Platforms, Inc., "Graph API Reference," Meta for Developers, 2023.
- [3] A. M. Kaplan and M. Haenlein, "Users of the world, unite! The challenges and opportunities of social media," *Business Horizons*, vol. 53, no. 1, pp. 59-68, 2010.
- [4] W. He, S. Zha, and L. Li, "Social media competitive analysis and text mining: A case study in the pizza industry," *Int. J. Inf. Manage.*, vol. 33, no. 3, pp. 464-472, 2013.
- [5] A. Joshi, P. Bhattacharyya, and M. J. Carman, "Automatic sarcasm detection: A survey," *ACM Comput. Surv.*, vol. 50, no. 5, pp. 1-22, 2017.
- [6] J. Eisenstein, "What to do about bad language on the internet," in *Proc. NAACL-HLT*, 2013, pp. 359-369.
- [7] H. Xu, M.-W. Chang, and C. Jayne, "A systematic review of social media-based sentiment analysis: Emerging trends and challenges," *Decis. Anal. J.*, vol. 3, 100073, 2022.
- [8] M. Rodríguez-Ibáñez, A. Casáñez-Ventura, F. Castejón-Mateos, and P.-M. Cuenca-Jiménez, "A review on sentiment analysis from social media platforms," *Expert Syst. Appl.*, vol. 223, 119862, 2023.
- [9] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. NAACL-HLT*, 2019, pp. 4171-4186.
- [10] J. Kreps, N. Narkhede, and J. Rao, "Kafka: A distributed messaging system for log processing," in *Proc. NetDB Workshop*, 2011, pp. 1-7.
- [11] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," in *Proc. ICLR*, 2013.
- [12] J. Pennington, R. Socher, and C. D. Manning, "GloVe: Global vectors for word representation," in *Proc. EMNLP*, 2014, pp. 1532-1543.
- [13] B. Pang, L. Lee, and S. Vaithyanathan, "Thumbs up? Sentiment classification using machine learning techniques," in *Proc. EMNLP*, 2002, pp. 79-86.
- [14] T. Joachims, "Text categorization with support vector machines," in *Proc. ECML*, 1998, pp. 137-142.
- [15] S. Rosenthal, N. Farra, and P. Nakov, "SemEval-2017 Task 4: Sentiment analysis in Twitter," in *Proc. SemEval*, 2017, pp. 502-518.

	Ms. Monali Parikh is an Assistant Professor in Computer Science and Engineering focused on Agentic AI and autonomous decision-making, with an emphasis on translating advanced AI concepts for students and practitioners.		Ms. Vaibhavi Pandya is an Assistant Professor in Computer Science and Engineering and current research interest includes machine learning and algorithms inspired by nature also having more than 5+ year experience in academics.
	Ms. Priyanka Makwana is an Assistant Professor in Computer Science academician with 8+ years of experience and areas of interest include Python, AI/ML Deep Learning, Network Security, DBMS, Software Engineering, and Computer Networks.		Ms. Barot Kajal is a Computer Science academician with 5+ years of experience and areas of interest include AI/ML, Data science, Cyber Security, Computer Networks, and JAVA Programming.