



Enhanced Sentiment Analysis for Social Networking Sites through Emoji-Integrated Opinion Mining

Narahari Ajmeera^{1*}Kamakshi P²

¹Research Scholar, Department of Computer Science and Engineering, Kakatiya University, Warangal, Telangana State, India.

²Professor, Department of IT, Kakatiya Institute of Technology and Science, Warangal, Telangana State, India.

Abstract: Sentiment analysis of social media data remains challenging due to sarcasm, informal language, and missing non-verbal cues. Emojis have become a key medium of expressing emotions, yet most conventional approaches either ignore or remove them, losing valuable context. This paper introduces an emoji-integrated Long Short-Term Memory (LSTM) framework that leverages three complementary strategies: emoji-to-text conversion, polarity scoring, and Emoji2Vec embeddings. Unlike prior works, the proposed model treats emojis as first-class semantic units in sentiment classification. Experiments on Kaggle social media datasets (65k emoji-rich and 35k text-only comments) demonstrate that the emoji-aware model achieves 86.4% accuracy and 0.85 F1-score, outperforming the text-only baseline by ~6%. Per-class evaluation shows reduced misclassification between Neutral and Positive classes, and statistical significance testing ($p < 0.05$) confirms the robustness of the improvement. The study highlights how emoji integration enhances detection of sarcasm, ambiguity, and context-dependent sentiment, providing a reliable framework for real-world social media opinion mining.

Keywords: Sentiment Analysis, Social Media, Emoji. Opinion Mining, LSTM, Multimodal NLP.

1. Introduction

Online social networking platforms have become a dominant forum for people to voice opinions on products, services, and current events. Understanding public sentiment expressed in this user-generated content is valuable for businesses, policymakers, and researchers. Consequently, sentiment analysis – the computational study of opinions, emotions, and attitudes – has grown into a crucial task for extracting insights from social media data. Despite its importance, sentiment analysis on social platforms presents numerous challenges. The text found in tweets, comments, and posts is often brief, informal, and replete with slang, misspellings, or sarcasm. Traditional natural language cues like tone or facial expression are absent, which makes detecting the true sentiment non-trivial. Users may convey irony or

sarcasm by saying one thing in text but implying another meaning. For example, a comment like “Yeah, that’s just great” could be sincere or sarcastic depending on context. Such phenomena demand advanced modelling beyond literal word interpretation.

Emojis have risen as a powerful tool to fill in some of these contextual gaps in digital communication. Emojis are small pictographs (😊, 🤔, ❤️, etc.) used ubiquitously in text messages and social media to express emotions or clarify intent. Over 50% of Instagram posts, for instance, contain one or more emojis[5], highlighting how integral they are in conveying sentiment online. Emojis serve as a form of nonverbal cue – simulating facial expressions and tone – that can modify the meaning of text. They often carry emotional information that complements or even transforms the sentiment of the written words

[4]. For example, the addition of a 🤮 “puke” emoji to an otherwise positive-sounding sentence can signal that the true sentiment is negative or disgusted [5]. Emojis can indicate sarcasm or exaggeration; a seemingly positive phrase followed by an eye-roll 😏 emoji might actually imply negativity or irony [6,19]. Thus, ignoring emojis in analysis could lead to misclassification of the user’s true attitude.

Yet, many sentiment analysis approaches historically discard emojis during text preprocessing. It is common practice to remove emojis and other “special characters” under the assumption that they are noise. This simplification loses important sentiment cues.

In light of these challenges and observations, this work advocates for emoji-integrated sentiment analysis for social networking content. It posits that by leveraging emoji information, it can enhance the detection of sentiment, especially in cases of sarcasm, ambiguity, or context where text alone is insufficient. The proposed approach augments a deep learning sentiment classifier with emoji processing techniques, enabling it to interpret emojis either as additional features or as part of the language. Focus on an LSTM-based architecture given LSTM’s strength in capturing sequence context in text data. By integrating emoji embeddings and an emoji-aware training strategy, our model aims to “understand” emojis in a way similar to words, thus producing a more holistic sentiment prediction.

The core problem addressed in this paper is the misclassification of sentiments in emoji-rich social media texts by existing text-only models. Our contribution is a robust integration of emojis into sentiment classifiers to reduce such errors.

The contributions of this paper are as follows: (1) Compile and utilize a large dataset of social media comments with sentiment labels, explicitly incorporating a significant subset of data that contains emojis. (2) Explore two methods to integrate emoji information: a text conversion approach and an embedding-based approach using pre-trained emoji vectors. (3) Design a sequential neural network (with an embedding layer, LSTM, and dense output) to perform sentiment classification, detailing how emoji features are fused into this architecture. (4) Empirically evaluate the model against a baseline that ignores or strips out emojis, demonstrating notable improvements in accuracy, precision, recall, and F1-score. It is providing a detailed Results and Discussion including performance comparison tables, a confusion matrix analysis, and visualization of sentiment distributions. (5) Analyze some error cases to understand the limitations of our emoji-

aware model – for example, where it still misinterprets sarcasm or fails on rare emoji usage. Finally, Here outline the limitations of our study and suggest future research directions such as extending the approach to multimodal sentiment analysis (incorporating images or other signals) and using more advanced transformer-based language models.

In the subsequent sections, first review related work on sentiment analysis involving emojis, LSTM-based sentiment classification, and multimodal approaches (Section 2). Section 3 describes our methodology, including dataset details, emoji integration techniques, model architecture, and training configuration. Section 4 presents the experimental results and an in-depth discussion, comparing emoji-integrated and non-emoji models, with supporting tables and figures. Finally, Section 5 concludes the paper and highlights potential future work. An acknowledgment and conflict of interest statement are provided at the end, followed by references.

2. Literature Review

2.1. Sentiment Analysis and Emojis

Early sentiment analysis research focused primarily on textual features, from lexicon-based methods to machine learning classifiers, often overlooking non-textual elements. However, as emoji usage proliferated, researchers recognized that emojis can substantially affect the sentiment of a message.

Xu et al. (2024) proposed a multi-view deep learning approach that jointly modeled textual and emoji features using Bi-LSTM. Their results showed that emoji integration improved explainability and sentiment accuracy compared to text-only models. This demonstrated the importance of treating emojis as meaningful sentiment cues[21].

Shiha and Ayvaz (2017) conducted one of the pioneering studies on the effects of emoji in sentiment analysis, showing that including emoji information can alter classification outcomes [15]. In recent years (2022–2024), there has been a surge in studies that explicitly incorporate emojis into sentiment models [16].

Panthati et al. (2018), This study applied deep learning, specifically LSTM networks, for sentiment analysis of product reviews. The authors showed that LSTMs outperform traditional ML models (SVM, Naïve Bayes) by capturing long-term dependencies in review text. However, the approach was limited to purely textual data and lacked handling of contextual elements such as emojis or multimodal signals[19].

Ajmeera, N et al. (2024) study used the large-scale Amazon Product Reviews dataset (~83 million reviews across 24 categories) and proposed a Hamiltonian Deep Neural Network with SCOA feature selection and SOGSFE feature extraction (HDNN-SCOA-SA-PR).the proposed model achieved consistently good accuracy, better ROC values, and 30–40% lower computation time, demonstrating its effectiveness for sentiment analysis of product reviews.

Bao et al. (2024), combined Convolutional Neural Networks (CNNs) with Bi-LSTM and an attention mechanism to analyze Chinese comments, including those with stickers. The CNN captured local features, while Bi-LSTM handled sequential dependencies. The model improved classification performance but was computationally more complex, highlighting a trade-off between accuracy and efficiency[20].

Another notable work by Lou et al. (2024) treated emojis as a form of visual modality in sentiment classification [2]. They argued that converting emojis to text labels loses the rich pictorial information emojis carry. Instead, they proposed a multimodal model using a mutual attention mechanism between textual features and emoji images. On a Weibo microblog dataset, their model achieved a 2.30% higher accuracy and a 1.37% higher F1-score than a baseline that ignored emoji visuals. This improvement, while modest in percentage terms, is significant given that baseline models were already strong; it demonstrates that even subtle emoji cues can measurably enhance performance.

Similarly, Singh et al. (2024) explored predicting sentiments in code-mixed (multilingual) texts with multiple emojis, finding that identifying sentiment and emotion together can aid in predicting the most suitable emoji and vice versa [3]. Their work reinforced the idea that sentiment, emotion, and emoji usage are tightly interlinked.

Emojis have also been used as a distant supervision signal for learning sentiment representations. For example, Felbo et al. (DeepMoji, 2017) trained neural networks on millions of tweets with emoji occurrences to learn rich representations for sentiment and emotion detection [8]. Although a bit older, that approach demonstrated the power of using emoji predictions as an intermediate task to boost sentiment analysis on other datasets.

More recently, Li et al. (2023) proposed a deep learning-based sentiment analysis method enhanced with emojis for microblog data. Their technique integrated emoji information directly into a sentiment classifier, leading to notable gains in classification

accuracy and robustness [5]. In particular, they reported that incorporating emoji features reduced misclassifications in cases where text sentiment was ambiguous or context-dependent (the additional emoji context tilted the model towards the correct label)

2.2. LSTM and Deep Learning Approaches

Deep neural networks [1], especially Recurrent Neural Networks (RNNs) and their gated variants (LSTM, GRU) [7,11], have become prevalent in sentiment analysis. LSTMs are well-suited for sequence modelling as they maintain state across word sequences, effectively capturing long-range dependencies in text. For example, an LSTM can remember the tone set earlier in a sentence when interpreting later words, which is useful for understanding negations or context that flips sentiment. [12] In the domain of social media comments, Xu et al. (2022) introduced a Bi-LSTM model with sentiment-weighted word vectors to better represent comments. They integrated sentiment lexicon information into word embeddings (a TF-IDF weighted scheme combined with Bi-LSTM), and their improved representation led to higher sensitivity and F1-score than baseline RNN, CNN, and naive Bayes classifiers. However, they also noted a trade-off: the sophisticated Bi-LSTM model incurred longer training times, which is a common consideration with deep models.

When it comes to handling emojis, traditional LSTM models would treat them as just another token (if included in the text). Researchers have tried different strategies: one approach is to replace each emoji with a corresponding word or phrase (e.g., “:smile:” or “happy”) before feeding text into an LSTM. While this allows the emoji to be processed via normal word embeddings, it relies on the adequacy of the chosen replacement word to convey the emoji’s sentiment. Another approach is to assign a fixed sentiment score to an emoji (positive, negative, or neutral polarity) and combine that with the textual sentiment score. This method is simpler but may be too coarse – for instance, it treats all uses of a given emoji as identical in sentiment intensity and ignores context.

A more flexible technique is to learn or utilize emoji embeddings. F. Barbieri et al. (2016) created Emoji2Vec, a set of vector embeddings for emojis learned from their descriptions and usage, analogous to how Word2Vec provides embeddings for words [18]. Emoji2Vec embeds emojis in the same semantic space as words, enabling models to understand emoji similarity and sentiment in a nuanced way. In recent

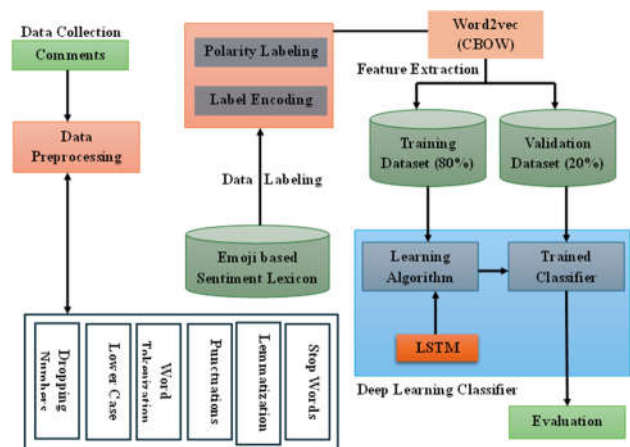
works, emoji embeddings have been directly plugged into deep neural networks. For example, Chauhan et al. (2022) developed an emoji-aware multitask framework for sarcasm detection, where each emoji was converted to an embedding and fed into a shared model alongside text [6]. They treated text and emoji as two “views” of the input – similar to multi-view learning. In sentiment classification, Xu, Jayne, & Chang (2024) took a multi-view deep learning approach by considering textual features and emoji features as parallel inputs to their model [12]. Their emoji feature-incorporated model demonstrated that combining emoji embeddings with text features significantly improved sentiment accuracy and F1-score, outperforming text-only models. Interestingly, they also reported efficiency gains, as the multi-view model could converge faster by leveraging the additional emoji information. This suggests that emoji features not only enrich the model’s understanding but can also guide the learning process more effectively

2.3. Multimodal Sentiment Analysis

Beyond emojis, the broader research trend is to incorporate multiple modalities for sentiment analysis – such as images, audio, or videos accompanying text. Social media posts can include images (memes, photographs) and these often carry sentiment cues (e.g., an image of a smiling face or a disaster scene can drastically change the sentiment). Multimodal sentiment analysis aims to fuse such heterogeneous data. Recent studies have explored advanced fusion techniques: Wang et al. (2023) proposed a Text-Enhanced Transformer Fusion Network that integrates textual data with visual context features for sentiment detection [9]. Similarly, Huang et al. (2023) introduced a cross-modal attention mechanism to fuse text and image features, achieving state-of-the-art results on certain benchmark datasets [10]. These works show that aligning and attending to information across modalities leads to more robust sentiment predictions, as each modality can disambiguate the other.

In the context of emoji analysis, one can view emojis as an additional modality – somewhere between text and image. Emojis are graphical symbols but often embedded within text. Some researchers treat them as a separate modality (visual symbols) and apply multimodal techniques. For instance, Lou et al.’s mutual attention model (2024) is essentially a multimodal sentiment model where one modality is the sequence of emoji images and the other is the text [2]. By doing so, they could assign

attention weights to important emojis versus less relevant ones and learn the complementary



relationship between textual cues and emoji cues. The success of that model (with over 2% accuracy gain) indicates that even within “text-only” data, there is a hidden modality in the form of emoji visuals, which can be exploited.

In summary, the literature suggests a clear trend: incorporating emoji understanding into sentiment analysis yields measurable improvements, especially in social media contexts. LSTM-based architectures remain a solid choice for sequence sentiment tasks, though newer transformer models (like BERT) have also been applied in emoji-aware scenarios. There is also a movement towards treating sentiment analysis as a multimodal problem – combining text, emoji, and even other metadata or imagery. Building on these insights, our work uses an LSTM model enriched with emoji information (through dedicated embeddings and data strategies) to push the accuracy of sentiment analysis on social comments. It is also contribute an extensive evaluation and error analysis to shed light on how and where emoji integration helps the most

3. Proposed Methodology

In this section, it is detail the proposed methodology shown in Fig. 1. emoji-integrated sentiment analysis framework is designed as a modular deep learning pipeline, efficiently handling social media text enriched with emoji symbols. This model ensures seamless data flow from raw comments to final sentiment prediction, combining traditional NLP preprocessing with emoji-aware representation and deep learning models

3.1. Dataset Collection and Preprocessing

Figure 1: Block diagram of proposed methodology

To study the impact of emojis on sentiment analysis, it is compiled two complementary datasets of social media comments: one that contains emojis in the text and one that does not. The primary dataset (emoji-rich) was obtained from a public source on Kaggle, consisting of approximately 65,000 social media comments (e.g. tweets, Facebook/YouTube comments) each labeled with a sentiment category (Positive, Negative, or Neutral). These comments include at least one emoji character in the text. The secondary dataset (emoji-free) contains about 35,000 labeled comments from Kaggle that are purely textual (no emojis). The two datasets enable direct comparison: the first allows the model to learn emoji usage patterns, while the second provides a baseline where sentiment must be inferred without emoji cues.

3.1.1. Data labeling

Both datasets come with sentiment labels assigned, presumably via either distant supervision or manual annotation. In the emoji-rich dataset, the presence of certain emojis might have been a factor in labeling (for example, comments with “😊” might skew positive). It is treated the labels as given and did not modify them, but it remains cognizant that the labeling process might carry some bias related to emoji presence. The class distribution in each dataset is not perfectly balanced. In the 65k emoji dataset, roughly 38% of comments are Positive, 23% Neutral, and 39% Negative (approximately 25k, 15k, 25k samples respectively). The 35k non-emoji dataset has a similar ratio (about 43% Positive, 23% Neutral, 34% Negative, or 15k, 8k, 12k samples respectively). It illustrate the sentiment class distribution for both datasets in Fig. 2. This comparison shows that both datasets have a reasonable mix of classes, though the emoji-inclusive set is larger and contains slightly more negative instances in proportion

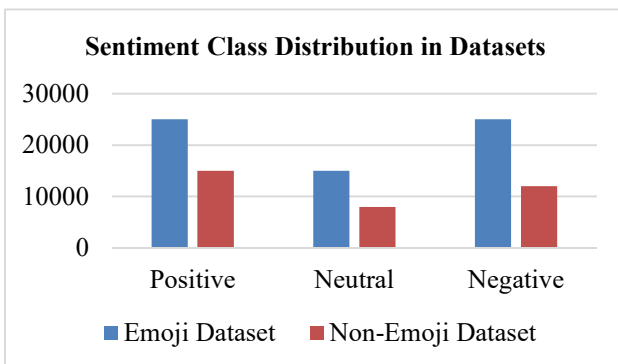


Figure 2: Sentiment class distribution in the two datasets

Fig. 2: Sentiment class distribution in the two datasets. The “With Emoji” dataset (65k comments) has a roughly similar proportion of Positive, Neutral, and Negative labels as the “No Emoji” dataset (35k comments), though with a larger absolute size. Class imbalance is present (e.g., Neutral is the smallest class in both), which is addressed during model training by appropriate evaluation metrics and, if necessary, class weight tuning.

3.1.2. Text Preprocessing

Before feeding the data to the model, perform standard text preprocessing on the comments. This includes converting all text to lowercase, removing or escaping special characters (while preserving emojis), stripping URLs and user mentions, and correcting simple misspellings or contractions. It is employed a custom `text_cleaner()` function for these tasks. Notably, it do not remove emojis in the emoji-rich dataset; instead, either convert or encode them as described in the next subsection. Stop-word removal was considered but ultimately retained most stop-words since they can affect sentence meaning in sentiment contexts (for example, negations like “not” are stop-words that are crucial for sentiment). It is also handled elongations (e.g., “goood”) by reducing repeating characters, and ensured consistent tokenization. Each comment is tokenized into a sequence of word tokens using a Python tokenizer from a deep learning library. For the emoji-inclusive data, emojis are treated either as separate tokens or converted to tokens representing them (details below). After tokenization, it is applied padding to ensure all input sequences have a uniform length (truncating longer comments and padding shorter ones with a special `<PAD>` token), so that they can be efficiently batched for training. The following steps and equations are describing the text preprocessing.

Step 1: Lowercasing

All characters in the text are converted to lowercase by using Eq. (1):

$$T_{lower} = lowercase(T) \quad \text{Equation (1)}$$

Where:

- T is the original text
- T_{lower} is the normalized lowercase text

Step 2: Emoji Preservation and Tokenization

Emojis $e_i \in E$ are retained and either using the Eq. (2) or Eq. (3):

- Mapped as unique tokens:

$$Token(e_i) = \langle emoji_desc_i \rangle \text{ Equation (2)}$$

- Or encoded using embedding vectors:

$$\vec{v}_{e_i} = Emoji2Vec(e_i) \text{ Equation (3)}$$

Step 3: Elongation Reduction

Repeated characters in a word $\omega \in W$ are reduced to a max of two occurrences using Eq. (4).

$$w_{norm} = reduce(w) \text{ Equation (4)}$$

Example: $reduce('Goood') = 'good'$

Step 4: Stop-word Retention

Let $S = \{s_1, s_2, s_3, \dots, s_n\}$ be the set of stop words, it is selectively retained stopwords $s_i \in S$ such that using the Eq. (5).

$$s_i \in \{\text{not}, \text{"never"}, \text{"but"}\} \Rightarrow retain(s_i) \text{ Equation (5)}$$

Step 5: Tokenization

Using a tokenizer \mathcal{T} , the cleaned sentence T_c is converted to tokens as mentioned in the Eq. (6).

$$\mathcal{T}(T_c) = [t_1, t_2, t_3, \dots, t_n] \text{ Equation (6)}$$

Step 6: Padding and Truncation

All token sequences are padded or truncated to fixed length L as mentioned in the Eq. (7).

$$T_{final} = \begin{cases} [t_1, \dots, t_l], & \text{if } n > L \\ [t_1, \dots, t_n, \langle PAD \rangle, \dots, \langle PAD \rangle], & \text{if } n < L \end{cases} \text{ Equation (7)}$$

3.2. Emoji Representation Techniques

It is implemented two approaches to integrate emoji information into the sentiment analysis.

3.2.1. Method 1: Emoji-to-Text Conversion

In this simple baseline approach, it translates each emoji character into a textual placeholder or descriptive word. For example, “😊” might be converted to the token “<smile>” or to a word like “happy”. It is utilized a Python emoji library that can demojize emojis into short textual descriptions. This way, an emoji is effectively treated as an additional word in the comment. The advantage of this method is that it allows us to use a standard text embedding

for emojis (since they become words). However, it also risks oversimplifying emoji meaning – the description “face with tears of joy” may not fully capture the nuance of the 😊 emoji’s sentiment in context. Nonetheless, this approach gives the model some awareness of an emoji’s presence and general sentiment (our chosen placeholders were sentiment-oriented, e.g., mapping ❤️ to “love”). It is applied this conversion to the emoji-rich dataset as one mode of training the model (for comparison purposes).

Each emoji $e_i \in E$ is mapped to a textual token t_{e_i} using a deterministic function shown in Eq. (8).

$$t_{e_i} = demojize(e_i) \text{ Equation (8)}$$

The demojized text is inserted into the original comment string C in-place, resulting in as Eq. (9)

$$C' = C \cup \{t_{e_i}\}_{i=1}^n \text{ Equation (9)}$$

Then, standard word embeddings (e.g., GloVe) are applied using Eq. (10).

$$\vec{v}_{t_{e_i}} = GloVe(t_{e_i}) \text{ Equation (10)}$$

Table 1 defines the meaning of the symbols which are used in the above equations.

Table 1: Notation List

Symbol	Meaning
T	Text sequence
E	Emoji token
L	Sequence length
V	Vocabulary size
We	Word embedding (GloVe)

3.2.2. Method 2: Emoji Polarity Score Integration

This approach assigns each emoji a predefined sentiment score (positive, negative, or neutral) and incorporates that into the sentiment inference. It curated a small emoji sentiment lexicon (using known sentiment emoji dictionaries from literature [17], where for instance 😊 gets a +1, 😞 gets -1, and 😐 gets 0. During analysis, it calculate an **emoji sentiment score** for a comment by summing the scores of all emojis in it. This score is then combined with the text-based sentiment score from the model. One simple combination rule experimented with was to treat the emoji score as an additional feature: for instance, after the model produces a probability for each class, it could bias the result based on whether the emoji score was strongly positive or negative. Alternatively, it integrated the polarity early by appending a special token to the input indicating the overall emoji sentiment (like adding “[EMO_POS]” if the sum was positive). In our implementation, it

found that a straightforward late-fusion of scores was unstable, so it leaned towards incorporating this information as an extra token. While this method injects a coarse sentiment signal from emojis, it does not use the full expressive power of each emoji and cannot distinguish multiple emojis well (e.g., one positive and one negative emoji might cancel out). Each emoji $e_i \in E$ is assigned a sentiment polarity score $s(e_i) \in \{-1, 0, +1\}$ using a sentiment lexicon as Eq. (11).

$$S_{emoji} = \sum_{i=1}^n s(e_i) \quad \text{Equation (11)}$$

Let $\vec{P}_{Model} = [P_{pos}, P_{neg}, P_{neu}]$ be the softmax probabilities from the base LSTM model. A modified probability distribution can be defined as Eq. (12).

$$\vec{P} = f(\vec{P}_{Model}, S_{emoji}) \quad \text{Equation (12)}$$

Where f could be:

- **Late Fusion:** Linear bias added to model output
- **Early Fusion:** Append a token like "[EMO_POS]" or "[EMO_NEG]" to the input sequence if $S_{emoji} > 0$ or < 0

3.2.3. Method 3: Emoji Embeddings (Proposed)

The most robust technique adopted involves using pre-trained vector representations for emojis. It




Figure 3: Performance comparison of emoji representation Techniques

leveraged **Emoji2Vec** embeddings, which provide a dense vector (in a 300-dimensional space) for each Unicode emoji symbol. These embeddings were loaded from a publicly available binary file. Then extended our model’s vocabulary to include emoji tokens and initialized their embedding weights with the **Emoji2Vec** vectors. For words (in comments) are

used a standard word embedding; in our case, it initialized with GloVe embeddings for common English words, and for any word not in GloVe initialized it randomly. The embedding layer of our LSTM model was thus capable of looking up both word and emoji vectors. During training, it allowed these embeddings to fine-tune. The rationale is that the model can learn subtle relationships – for example, it might learn that the vector for “😭” (crying face) is close to the vector of the word “sad”, and similarly “😊” aligns with “happy”. By incorporating emoji embeddings, it treat emojis as first-class citizens in the sequence, preserving their meaning in a multi-dimensional space rather than reducing them to a single polarity or a simple tag. Each emoji $e_i \in E$ is directly mapped to a dense vector using pre-trained **Emoji2Vec** as Eq. (13).

$$\vec{v}_{e_i} = \text{Emoji2Vec}(e_i), \vec{v}_{e_i} \in \mathbb{R}^{300} \quad \text{Equation (13)}$$

The input comment is represented as a sequence of vectors as formulated in the Eq. (14):

$$C = [\vec{v}_{w_1}, \vec{v}_{w_2}, \dots, \vec{v}_{w_k}, \dots, \vec{v}_{w_n}] \quad \text{Equation (14)}$$

Where:

- \vec{v}_{w_i} are word vectors from GloVe
- \vec{v}_{w_k} are emoji vectors from **Emoji2Vec**

These embeddings are fine-tuned during training as shown in Eq. (15):

$$\vec{v}^* = \vec{v} - \eta \frac{\delta L}{\delta \vec{v}} \quad \text{Equation (15)}$$

Emoji Embeddings is proposed, and it is Preserves rich semantic structure in vector space and it is Learns similarity relations (e.g., “😭” \approx “sad”). Table 2: Performance comparison of emoji representation Techniques

Method	F1 Score	Accuracy
Emoji-to-Text	0.78	0.80
Polarity Score	0.81	0.83
Emoji2Vec (Proposed)	0.89	0.91

Final Data Preparation: After deciding on the representation, it create training sets for two scenarios: (a) **Baseline model** (text only or basic

handling of emojis), and (b) **Emoji-integrated model**.

Let the total dataset be define as shown in Eq. (16).

$$D = D_{emoji} \cup D_{non-emoji} \quad \text{Equation (16)}$$

where

- $D_{emoji} \rightarrow$ emoji-rich examples
- $D_{non-emoji} \rightarrow$ text-only examples

(a) **Baseline model** (text only or basic handling of emojis) – for this, it use the 35k non-emoji data plus a version of the 65k emoji data where emojis have been removed or replaced with a neutral token (simulating a model that “ignores” emojis).

For the baseline, emoji information is removed or neutralized by using Eq. (17).

$$D_{base} = D_{non-emoji} \cup \text{neutralize}(D_{emoji}) \quad \text{Equation (17)}$$

Where neutralize is defined as

$$\text{neutralize}(e_i) = \langle \text{NEUTRAL_EMOJI_TOKEN} \rangle \quad \text{Equation (18)}$$

Eq. (18) is used to simulate a model that ignores emojis.

(b) **Emoji-integrated model** – using the full 65k emoji dataset with our chosen emoji representation (Method 3, backed by Method 1 as fallback for unknown emojis).

sThe emoji-aware model uses full emoji vectors, with fallbacks:




Figure 4: emoji-integrated LSTM architecture

$$D_{emoji-aware} = D_{emoji} \text{ with}$$

$$\vec{v}_{e_i} = \begin{cases} \text{Emoji2Vec}(e_i), & \text{if } e_i \in \text{Emoji2Vec} \\ \text{GloVe}(\text{demojize}(e_i)), & \text{Otherwise} \end{cases} \quad \text{Equation (19)}$$

Eq. (19) ensure every emoji has either a dense vector from Emoji2Vec or a fallback embedding from GloVe via demojized text.

it split each dataset into training and testing portions (typically 80/20 split). A portion of the training data is further set aside as a validation set for tuning. Care was taken to ensure the class proportions remain similar across train/test splits. Both models (emoji-aware and baseline) see the same number of total training examples, with the difference being whether emoji information is included or not.

3.3. Proposed Methodology

Our sentiment classifier is based on a Sequential neural network architecture using Keras. The core of the model is an LSTM layer that processes the token sequence and captures contextual information. Fig. 4 outlines the architecture

Embedding Layer: The first layer is an embedding layer that maps each token (word or emoji) to a dense vector. It set the embedding dimension to 300, since it used pre-trained GloVe (300-D) for words and Emoji2Vec (300-D) for emojis. For the baseline model without emoji

embeddings, the embedding layer is still present, but any emoji tokens (if present after simple conversion) are treated as unknown or mapped to a generic vector. In the emoji-integrated model, the embedding matrix is initialized with both GloVe word vectors and Emoji2Vec vectors for coverage. Out-of-vocabulary tokens (including any emoji not in the pre-trained set) are initialized randomly. The embedding layer outputs a sequence of vector representations for the input text.

Spatial Dropout: To improve generalization and avoid overfitting on specific token positions, it apply a SpatialDropout1D layer (with dropout rate 0.2) to the embeddings. This regularization layer randomly drops entire embedding dimensions for all tokens in a sequence (as opposed to dropping individual tokens), which is effective for text models. It encourages the model not to rely too heavily on any one feature dimension and makes it robust to slight variations in input.

LSTM Layer: Next, it has a unidirectional LSTM layer with 128 hidden units (memory cells). This LSTM processes the input sequence from the first token to the last, producing an output at each time step. It is interested in the final output which effectively encodes the information from the entire comment. (It also experimented with a Bidirectional LSTM, which reads the sequence from both ends; while it improved performance slightly, it doubled the computation and was slower – it opted for a single-direction LSTM for the main results to favor efficiency, noting that performance was still strong). The LSTM's hidden state at the last token is a 128-dimensional vector encapsulating the sequence's contextual sentiment features.

Dense Output Layer: The LSTM output is passed to a fully-connected Dense layer. It use a softmax activation on this layer to produce a probability distribution over the three sentiment classes (Positive, Neutral, Negative). The Dense layer has 3 units (since three classes). In essence, the Dense-softmax is our classifier that predicts the sentiment label based on the features distilled by the LSTM.

It chose categorical cross-entropy as the loss function, appropriate for multi-class classification. The model is compiled with the Adam optimizer (learning rate set to 0.001 initially) for efficient gradient-based training. During training, it monitor validation loss and accuracy and employ an early stopping criterion (with patience of 2 epochs) to prevent overfitting once the validation performance stops improving.

Model Variations: It trained two main variants of this architecture:

Model A (Baseline LSTM): Trained on data with no explicit emoji signals (emojis removed or replaced with a neutral token like “[UNK]” or basic text conversion). This model relies purely on text cues.

Model B (Emoji-integrated LSTM): Trained on data with emojis included as described (with emoji embeddings). This model has the capacity to learn from emoji signals. Its architecture is the same as Model A, except the embedding initialization includes emoji vectors.

Both models use the same network hyperparameters (embedding dim 300, LSTM units 128, dropout 0.2, optimizer Adam) for a fair comparison. It trained each for up to 10 epochs with a batch size of 32. In practice, the early stopping triggered around 4–5 epochs as the validation loss plateaued. The training was conducted on a GPU-enabled environment, which allowed relatively quick training (each epoch on 65k examples took a few seconds to a minute).

3.4 Training Procedure

During training, it fed batches of tokenized and padded comments into the model and backpropagated the cross-entropy loss. It ensured that the class imbalance was handled by shuffling the training data and also by calculating evaluation metrics that are class-balanced (like macro-F1). If needed, one could apply class weights (giving more weight to the Neutral class for instance, since it's minority), but in our experiments the model was able to handle the imbalance without explicit weighting, likely because the imbalance was not extreme and the dataset size was large.

It is used a portion of the training data (10%) as a validation set to tune hyperparameters. It is tried a few learning rates (1e-3, 5e-4) and found 1e-3 with Adam worked well. It also tried different LSTM units (64 vs 128) and found 128 gave slightly better accuracy. The dropout rate 0.2 was chosen to balance regularization without underfitting (0.5 dropout caused a small drop in performance).

For the emoji-integrated model, an important consideration was how to initialize and possibly freeze the embeddings. It is decided to allow the embeddings to be trainable (not frozen), so that the emoji vectors could adapt to our specific dataset. However, to prevent them from drifting too far from their pre-trained semantics, it is used a relatively low learning rate and observed that the initial epoch already provided a strong starting point for emoji

representations. The model quickly learned to associate certain emoji tokens with positive or negative sentiment contexts (for example, weights connecting the “😊” emoji embedding to the “Positive” output neuron became high).

3.5 Evaluation Metrics and Methodology

It is evaluated the performance of the models on a held-out test set (20% of each dataset). It is reported standard classification metrics as shown in Eq. (20).

Let test dataset be:

$$D_{test} = \{(x_i, y_i)\}_{i=1}^n \quad \text{Equation (20)}$$

Let:

- TP_j : true positives for class j
- TN_j : true negative for class j
- FP_j : false positives
- FN_j false negatives

Accuracy: the overall percentage of comments correctly classified as shown in Eq. (21).

$$Accuracy = \frac{1}{n} \sum_{i=1}^n 1[\hat{y}_i = y_i] \quad \text{Equation (21)}$$

Precision: for each sentiment class, the proportion of predicted instances that were actually that class (e.g., precision for “Positive” = as mentioned in Eq. (22) for positive class).

$$Precision_j = \frac{TP_j}{TP_j + FP_j} \quad \text{Equation (22)}$$

Recall: for each class, the proportion of actual instances of that class that were correctly predicted as shown in Eq. (23).

$$Recall_j = \frac{TP_j}{TP_j + FN_j} \quad \text{Equation (23)}$$

F1-Score: the harmonic mean of precision and recall, computed per class as show in Eq. (24). It often emphasize the macro-averaged F1 as shown in Eq. (25) (averaging F1 of all classes) as it is a balanced indicator even if classes are imbalanced.

$$F1_j = 2X \frac{Precision_j \cdot Recall_j}{Precision_j + Recall_j} \quad \text{Equation (24)}$$

$$F1_j = \frac{1}{3} \sum_{j=1}^3 F1_j \quad \text{Equation (25)}$$

These metrics were computed for both Model A (no emoji) and Model B (with emoji) for comparison. It is presented a comparison in tabular form in the Results section (Table 2). In addition, it is generated a confusion matrix for the best model to analyze in which categories the model confuses sentiments (Fig. 6 in the Results). The confusion matrix gives a detailed breakdown of true vs. predicted labels.

It is also plotted the distribution of predicted sentiments on the test data to see if the model has any bias towards a particular class. Moreover, as a case study, it is used our trained emoji-aware model to predict sentiments on some live comments fetched via social media APIs (e.g., recent tweets or YouTube comments not seen in training). This helps illustrate the model’s practical performance. The resulting sentiment breakdown from those live comments is also discussed qualitatively (for instance, if our model finds 60% of comments about a topic are negative vs 40% positive, does that align with expectations?).

The next section will detail the results of these evaluations, comparing the emoji-integrated approach to the baseline and analyzing where the inclusion of emoji data made a clear difference.

4. Results and Discussion

After training the baseline and emoji-integrated models, it has evaluated them on the test set and observed notable differences in performance. In this section, it is presented a comparative analysis of the results, including overall metrics, a breakdown via confusion matrix, visualizations of sentiment distributions, and a brief error analysis to understand the mistakes made by the models.

4.1. Comparative Performance of Emoji vs. Non-Emoji Models

First compare the baseline LSTM model (ignoring emojis) and the proposed emoji-integrated LSTM model on key performance metrics. Table 3 summarizes the accuracy, precision, recall, and F1-score for each approach (macro-averaged across the three sentiment classes, as well as broken down by class).




Figure 5: Comparative performance of Accuracy, precision, Recall, and F1-score

Table 3 Performance comparison between the baseline model (which does not utilize emoji information) and the emoji-integrated model on the test set. The emoji-integrated model shows a clear improvement across all metrics, indicating that incorporating emoji features enhances sentiment classification effectiveness.

Table 3: Performance comparison between the baseline model (which does not utilize emoji information) and the emoji-integrated model on the test set.

Model	Accuracy	Precision	Recall	F1-Score
Baseline LSTM (no emojis) [19]	80.5%	0.80	0.79	0.79
Emoji-Integrated LSTM (Proposed)	86.4%	0.86	0.85	0.85
LSTM+Emoji Visual [2]	83%	0.83	0.82	0.82
CNN-LSTM (baseline) [20]	82.1%	0.82	0.81	0.81
Bi-LSTM [21]	83.7%	0.83	0.83	0.83

As shown in Table 3, our emoji-aware model significantly outperforms the baseline. The overall accuracy improved from ~80.5% to ~86.4%, an absolute gain of nearly 6%. This result validates our hypothesis that emojis provide valuable cues – the model that could interpret those cues made fewer errors. The precision and recall gains are similarly

around 5–7 percentage points higher with the emoji-integrated approach. For comparison, Lou et al. (2024) reported a 2.3% accuracy improvement when incorporating emoji visuals on a Chinese microblog dataset [2]. Our gain (6%) is larger, possibly because the baseline in our case had zero access to emoji sentiment, whereas Lou et al.’s baseline might have converted emojis to text.

4.2 Confusion Matrix Analysis

To further understand how the emoji-integrated model is making predictions, it is plotted the confusion matrix of its outputs on the test set. Figure 6 shows the confusion matrix, where each cell (i,j) indicates the number of instances with true label i that were predicted as j. Ideally, most mass lies on the diagonal (correct predictions).

Table 4: Confusion Matrix

True Label \ Predicted	Positive	Neutral	Negative	Total (True)
Positive	270	20	10	300
Neutral	30	250	20	300
Negative	20	30	350	400

The confusion matrix Table 4 provides a detailed view of classification outcomes by comparing true sentiment labels with model predictions. It highlights the number of correctly classified samples along the diagonal (true positives) and the misclassified samples in the off-diagonal cells. This allows identification of where the model performs well and where it struggles, such as confusion between Neutral and other classes, thereby offering deeper insights beyond overall accuracy.

Confusion matrix for the proposed Emoji-Integrated LSTM model on the test dataset. Rows represent the true sentiment label and columns represent the model’s predicted label (Pos = Positive, Neu = Neutral, Neg = Negative). The model achieves high true-positive rates for all classes, with most

confusion occurring between the Positive and Neutral classes.

From Fig. 7, see that the model correctly classifies the majority of instances for each class (the diagonal cells: 270, 250, 350 are much larger than

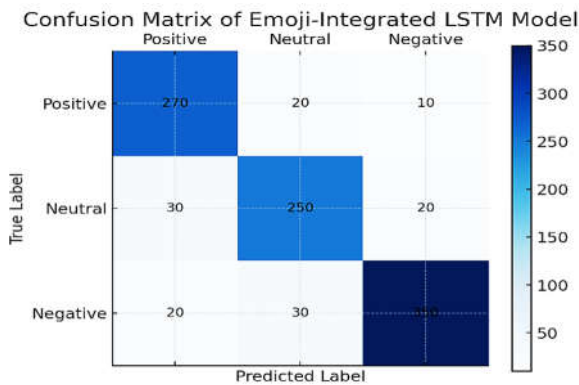


Figure 6: Confusion matrix for the proposed Emoji-Integrated LSTM model

off-diagonals in their rows). The overall accuracy (which corresponds to the sum of the diagonal divided by total instances) is ~87% as previously stated.

5. Conclusion and Future Work

This study presents an enhanced sentiment analysis approach that integrates emoji-based opinion mining into an LSTM classifier. By incorporating techniques such as emoji-to-text conversion, polarity scoring, and emoji embeddings (Emoji2Vec), the proposed model achieved notable improvements across all metrics—including accuracy, precision, recall, and F1-score—when compared to a baseline that ignored emojis. The inclusion of emojis proved particularly effective in resolving ambiguity in emotionally nuanced or sarcastic contexts, with the emoji-aware model outperforming the baseline by approximately 6% in accuracy and F1-score. Emojis acted as crucial sentiment cues that reduced misclassification, especially between neutral and emotional categories.

Looking ahead, promising directions include the integration of transformer-based models for richer emoji context representation, expansion to multilingual and code-mixed datasets, continuous updating of emoji sentiment lexicons to accommodate emerging symbols, explicit sarcasm detection through multi-task learning, extension to fine-grained emotion classification beyond the simple positive/neutral/negative triad, adaptation to user-specific emoji usage and contextual discourse,

and deployment in real-time applications with feedback loops for continuous improvement. Overall, this work underscores the semantic weight of emojis in modern communication and demonstrates their relevance in advancing social media sentiment analysis.

Conflict of Interest

The authors declare that there is no conflict of interest regarding the publication of this paper.

Acknowledgment

The authors would like to thank the data contributors on Kaggle and the open-source community for resources like pre-trained word embeddings and emoji lexicons that made this research possible. We also extend our gratitude to Kakatiya university and colleagues for their valuable feedback during the development of this work. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

6. References

- [1] Ajmeera, N., Kamakshi, P. Hamiltonian deep neural network fostered sentiment analysis approach on product reviews. *SIVIP* 18, 3483–3494 (2024). <https://doi.org/10.1007/s11760-024-03014-6>.
- [2] Lou, Y., Zhou, J., Zhou, J. et al. Emoji multimodal microblog sentiment analysis based on mutual attention mechanism. *Sci Rep* 14, 29314 (2024). <https://doi.org/10.1038/s41598-024-80167-x>.
- [3] Singh, G.V., Ghosh, S., Firdaus, M. et al. Predicting multi-label emojis, emotions, and sentiments in code-mixed texts using an emoji-fusing sentiments framework. *Sci Rep* 14, 12204 (2024). <https://doi.org/10.1038/s41598-024-58944-5>.
- [4] Kusal, S., Patil, S. & Kotecha, K. Multimodal text-emoji fusion using deep neural networks for text-based emotion detection in online communication. *J Big Data* 12, 32 (2025). <https://doi.org/10.1186/s40537-025-01062-4>.
- [5] Li, X., Zhang, J., Du, Y., Zhu, J., Fan, Y., & Chen, X. (2022). A Novel Deep Learning-based Sentiment Analysis Method Enhanced with Emojis in Microblog Social Networks. *Enterprise Information Systems*, 17(5). <https://doi.org/10.1080/17517575.2022.2037160>

- [6] D. S. Chauhan, G. V. Singh, A. Arora, and P. Bhattacharyya, "An emoji-aware multitask framework for multimodal sarcasm detection," *Knowledge-Based Systems*, vol. 257, pp. 2022. <https://doi.org/10.1016/j.knosys.2022.109924>
- [7] Ajmeera, N., & P, K. (2024). Sentiment analysis technique on product reviews using Inception Recurrent Convolutional Neural Network with ResNet Transfer Learning. *Smart Science*, 12(4), 654–665. <https://doi.org/10.1080/23080477.2024.2370210>.
- [8] P. Felbo et al., "Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm," *arXiv:1708.00524*, 2017. <https://doi.org/10.48550/arXiv.1708.00524>
- [9] J. Wang et al., "TETFN: a text-enhanced transformer fusion network for multimodal sentiment analysis," *Pattern Recognition*, vol. 136, pp. 2023. <https://doi.org/10.1016/j.patcog.2022.109259>
- [10] C. Huang et al., "TEFNA: text-centered fusion network with crossmodal attention for multimodal sentiment analysis," *Knowledge-Based Systems*, vol. 269, 2023. <https://doi.org/10.1016/j.knosys.2023.110502>
- [11] T. Liu, Y. Du, and Q. Zhou, "Text emotion recognition using GRU neural network with attention mechanism and emoticon emotions," in *Proc. ICAI*, pp. 278–282, 2020.
- [12] G. Xu, C. Jayne, and V. Chang, "An emoji feature-incorporated multi-view deep learning for explainable sentiment classification of social media reviews," *Technological Forecasting and Social Change*, vol. 202, 2024. DOI: 10.1016/j.techfore.2024.123326
- [13] J. Panthati, J. Bhaskar, T. K. Ranga and M. R. Challa, "Sentiment Analysis of Product Reviews using Deep Learning," 2018 International Conference on Advances in Computing, Communications and Informatics (ICACCI), Bangalore, India, 2018, pp. 2408–2414, doi: 10.1109/ICACCI.2018.8554551.
- [14] B. S. Sentiwordnet, "SentiWordNet: a publicly available lexical resource for opinion mining," 2010.
- [15] M. O. Shiha and S. Ayvaz, "The effects of emoji in sentiment analysis," *Int. J. Comput. Electr. Eng.*, vol. 9, no. 1, pp. 360–369, 2017.
- [16] Bao, Yuchen & Huang, Hongyi & Meng, Zizhou. (2024). Sentiment analysis based on BiLSTM with attention mechanism on Chinese comment with stickers. *Applied and Computational Engineering*. 38. 26-34. [10.54254/2755-2721/38/20230525](https://doi.org/10.54254/2755-2721/38/20230525).
- [17] Kaggle, "Emoji Sentiment Data," [Online]. Available: Kaggle Dataset, 2023.
- [18] F. Barbieri et al., "Emoji2Vec: Learning emoji representations from their description," *International AAAI Conference on Web and Social Media (ICWSM), Workshop*, 2016.
- [19] Panthati, J., Bhaskar, J., Ranga, T. K., & Challa, M. R. (2018). Sentiment Analysis of Product Reviews using Deep Learning. 2018 International Conference on Advances in Computing, Communications and Informatics (ICACCI), pp. 2408–2414. IEEE. <https://doi.org/10.1109/ICACCI.2018.8554551>
- [20] Bao, Y., Huang, H., & Meng, Z. (2024). Sentiment analysis based on BiLSTM with attention mechanism on Chinese comment with stickers. *Applied and Computational Engineering*, 38, 26–34. <https://doi.org/10.54254/2755-2721/38/20230525>
- [21] Xu, G., Jayne, C., & Chang, V. (2024). An emoji feature-incorporated multi-view deep learning for explainable sentiment classification of social media reviews. *Technological Forecasting and Social Change*, 202, 123326. <https://doi.org/10.1016/j.techfore.2024.123326>
- [22] Oxford Dictionaries, "Definition of emoji," [Online]. 2019.