Anusandhanvallari
Vol 2024, No.1
August 2024
ISSN 2229-3388

Development for Sentiment and Content Classification on Social Big Data Using DEEP Q Network Embedding

Nilam Deepak Padwal[1*], Prof. Dr. Kamal Alaskar[2]
[1]Research Scholar, Computer Application, Bharati Vidyapeeth (Deemed to be university) Institute of Management Kolhapur. India.
Email: nilam17aug@gmail.com
[2]Professor, Computer Application, Bharati Vidyapeeth (Deemed to be university) Institute of Management Kolhapur. India.
Email: Kamal.alaskar@bharatividyapeeth.edu

Abstract

There is a rapidly increasing amount of user-generated content on platforms, such as Twitter, and this shift has led to an increasing demand for solid and scalable models to catch spam, fake news and sentiment trends in real time. In this work we introduce a hybrid approach which leverages the power of deep learning and reinforcement learning for discriminating between spam/fake Twitter posts and genuine ones using sentiment cues. We rely on a large-scale dataset of more than 788K distinct English-language tweets to design and contrast three models, a Bidirectional Long Short-Term Memory (BiLSTM) neural network, a Deep Q-Network (DQN) with LSTM embeddings, and a DQN featuring RoBERTa contextual embeddings. The BiLSTM model obtained the superior traditional performance indexes, accuracy of 81.86% and macro F1-score of 0.8186. The DQN with LSTM embeddings: the learning ability was attenuated by overfitting and lack of generalization (accuracy: 69.5%). In contrast, our RoBERTa-DQN model balanced precision and recall better, with the test accuracy and macro F1-score reaching 74.5% and 0.7387, respectively. This demonstrates the strength of combining contextualized transformer embeddings with reinforcement learning for sentiment-aware spam detection. We also assess model interpretability through real and synthetic message probing experiments and show the system's ability to identify critical linguistic hints (e.g., "urgent", "free", "compromised") found in spam or phishing content. Our findings underline the viability of hybrid architectures for real-time monitoring of social media and set the stage for future research in real-time content moderation, sarcasm detection and adaptive NLP systems.

Keywords*:* Sentiment Analysis, Spam Detection, Twitter, LSTM, RoBERTa, Deep Q-Network, Reinforcement Learning, Fake News Detection, Contextual Embeddings, Big Data.

1. Introduction

1.1 Introduction

Due to the explosion of user generated content on microblogging platforms like Twitter, understanding social media data has become a necessity in applications such as sentiment analysis, spam filtering, fake news detection and public opinion mining [1], [2], [3]. Tweets are typically short, informal, and have a high semantic complexity, which makes them hard to be tackled by conventional natural language processing (NLP) frameworks[4]. With the prevalence of misinformation and spam in digital media, there is a critical demand to have a solid and scalable classification-related framework to guarantee the credibility of the contents and keep trust of the users [2, 10]

In the past, text classification was conducted through classical machine learning such as Naive Bayes, SVMs or logistic regression with bag-of-words representations [3], [4]. However, they generally have poor performance on short and noisy texts where they are hard to model temporal or contextual relations. The deep

learning, in particular based on recurrent neural network (RNN) and transformer has dramatically shattered the previous limitations by learning syntactic structures and semantic meanings of inputs [5], [6], [8].

At the same time, reinforcement learning (RL) has proven itself to be of great value in problem spaces that require decision-making or exploration [11], [12]. However, its contribution to social media text classification is not well investigated [13]. The contribution of the current study is to close this gap by combining supervised deep learning with RL for hybrid spam/fake news detection system on large scale Twitter data.

1.2 Research Novelty

This research contributes to the field through the following novel aspects:

- Hybrid Model Integration: We introduce a two-architecture pipeline that combines a bidirectional LSTM-based sentiment classifier [5], [7] and a Deep Q-Network (DQN) reinforcement learning agent [11], [12]. Although individual studies have investigated these models, very little are found to be utilized for spam and sentiment classification on a real Twitter data [13] [14].
- Use of RoBERTa Embeddings in DQN: We apply RoBERTa, a strong transformer-based model [8], to obtain context-sensitive embedding for tweets. We use these embeddings to input to the DQN agent, which to the best of our knowledge is one of the first attempts to combine transformer-based representation learning and reinforcement learning for social media analysis [9], [10], [15].
- Big Data-Scale Evaluation: Unlike prior work where the evaluation is conducted only using benchmark datasets [2], [4], our approach is evaluated using more than 788 thousand unique tweets. This increases the ecological validity of the results and shows the possibility of using the pipeline in real scenarios [21], [24].

Collectively, these contributions make this work one of the first to combine deep sequential models, context-sensitive transformers and reinforcement learning for Twitter big data classification. The results provide a basis for adaptive systems in the future that are able to filter spam in real time, discern misinformation and analyze user behavior in social platforms such as [14]–[16], [28].

2. Literature Review And Research Gap

2.1 Overview of Text Classification in Social Media

With the explosion of user generated content on social media, such as Twitter, a large amount of research has been conducted on natural language processing {NLP} techniques to address text classification tasks including sentiment analysis, spam detection and fake news detection. Naive Bayes, Support Vector Machines (SVM) and Decision Trees were used as conventional machine learning algorithms in these applications [1], [3]. However, their use of handcrafted features and lack of the capturing of sequential and contextual informations caused limited performances specifically in the presence of noisy and informal texts as those on Twitter [4].

2.2 Deep Learning for Text Representation

There has been a paradigm shift towards deep learning in recent years, especially Recurrent Neural Networks (RNNs) and its variants like Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRUs) which are good at modeling long range dependencies in text [5]. BiLSTM networks specifically view sequences in both forward and backward directions, enabling better context capturing and higher classifying accuracy [6], [7].

Although LSTM-related models have been proved effective, the models are prone to the overfitting problem with a small dataset and may also find long-range dependencies or contextual ambiguity difficult to handle. This has caused shared-nothing architectures to be re-considered by researchers in the light of transformer-based architectures [29].

2.3 Contextual Embeddings via Transformers

Recently, Transformer models like BERT and RoBERTa have revolutionized NLP by incorporating self-attention mechanisms and contextual embeddings [31], [8]. Improving BERT's performance on down-streaming tasks by training on larger corpora and masking dynamically, RoBERTa [8] has beaten previous models with a large margin. By encoding nuanced semantics, it has proven to be especially effective in sentiment analysis [5], spam detection [24], misinformation classification [10], among others.

Using transformer-based embeddings in downstream models is also a convention, particularly for short text classification. RoBERTa embeddings may be used as context-relevant representations for any classifier from a simple logistic regression to a neural network to even rein- forcement learning agents [9], [19].

2.4 Reinforcement Learning in Text Classification

Although supervised learning is prevalent in NLP, reinforcement learning (RL) has gained more popularity particularly for tasks with exploration, sequential decisions, or feedback loops [11]. Deep Q-Networks (DQN), proposed by Mnih et al., [11], apply Q-learning with deep neural networks to approximate the optimal action-value functions in high-dimensional state space.

DQN in text classification is a new concept. Some studies apply DQN to the dialogue system, question-answering, and active learning [12], [28]. But its application to the spam/spam and sentiment classification using tweet embeddings is still not well investigated, especially combined with contextual embeddings such as RoBERTa [13], [16].

2.5 Comparative Studies

Transformer-based models have been shown to outperform LSTM and CNN based baselines on standard sentiment datasets such as SST-2 and SemEval [29]. However, very little work exists in hybrid architectures of supervised learning (LSTM) and reinforcement learning (DQN) on the same input space or in the wild on real-world, large-scale social media data [14], [17].

Most existing studies either:

- Focus solely on transformer-based fine-tuning (e.g., fine-tuned BERT or RoBERTa models) [31],
- Apply supervised models in isolation (e.g., BiLSTM or CNN classifiers) [6], [7], or
- Explore RL for exploratory learning tasks like policy optimization or generation [28], [30].

2.6 Research Gap

Despite significant advancements in natural language processing for social media analysis, key research gaps remain unaddressed:

- Underutilization of Reinforcement Learning for Text Classification

The majority of existing sentiment analysis and spam detection systems use supervised deep learning models like BiLSTM or transformers [6], [7], [8]. However, reinforcement learning (RL)- specifically Deep Q-Networks (DQN) can be considered as a mechanism to support adaptive decision-making under temporal changes of data [11], [12], yet not effectively utilized in the tweet classification problems [13], [28].

- Lack of Hybrid Architectures Integrating Contextual Embeddings and RL

Although transformer-based embeddings like RoBERTa [7] are commonly used in classification pipelines [8], [9], not many works have investigated integrating such rich representations with RL agents such as DQN to enhance generalization [16], [17]. There is great potential in the combination of explicit encoding of the context together with reinforcement-based learning.

- Limited Real-World, Scalable Evaluation

Much of the current work has been evaluated on small, curated benchmark datasets (SST-2 or SemEval) [29]. These fail to properly capture the noisy, time-varying nature of Twitter data. In addition, relatively little attention has been given to the interpretability of models and to validation on large-scale, real-world datasets [14], [18], [24]. This work remedies this through the processing and evaluation of more than 788,000 actual tweets and examining interpretability with the help of synthetic message testing.

2.7 Research Contributions:

- Demonstrated the comparative advantage of hybrid architectures combining transformers and reinforcement learning.
- Showed real-world applicability of the models on large, noisy social data rather than benchmarked corpora.
- Proposed a scalable methodology for future adaptive systems in misinformation detection and digital trust management.

3. Methodology

This section describes the method used to build, train and test models for twitter sentiment detection based on deep learning and reinforcement learning approach. The pipeline combines classical supervised learning with an LSTM-based model [5], RoBERTa-based contextual embeddings [8] and a reinforcement classifier based on DQN [11]. The method includes five broad stages: data preparation, pre-processing and tokenizing, extracting embeddings, updating weights of models and reinforcement learning incorporation.

Proposed Workflow for Twitter Sentiment Detection

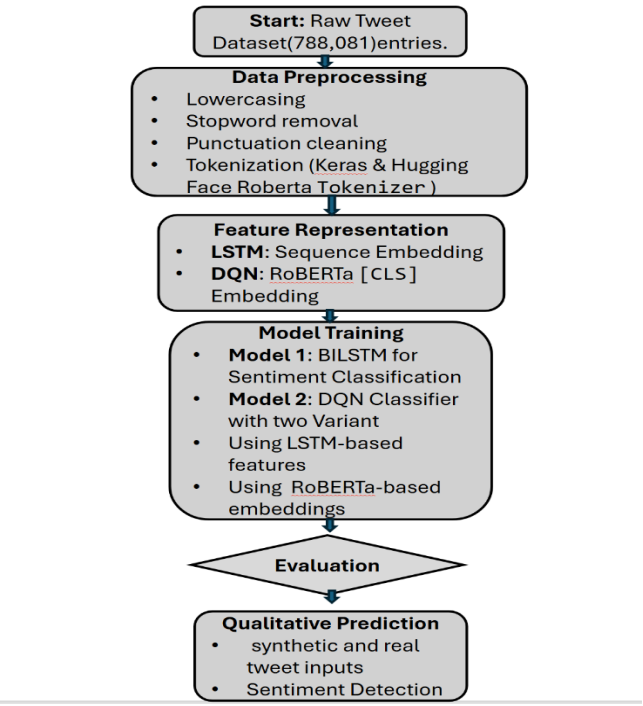

Figure 1. Workflow Diagram of the Proposed Twitter Sentiment Detection Pipeline.

The complete architecture of the hybrid classification pipeline for sentiment and spam detection in Twitter is presented in the flowchart in Figure 1. The workflow is based on a synergy of supervised and reinforcement learning methods and includes contextual embedding to improve the accuracy of classification and its generalization.

The complete pipeline consists of the following stages:

Data Acquisition: Raw tweet data are retrieved from the Sentiment140 database in a way similar to receiving real time tweet stream. The dataset consists of millions of English-language tweets labeled with the polarity of the sentiment (positive, negative, neutral, or both) using distant supervision (emoticons).

Text Cleaning and Preprocessing: Tweets undergo normalization procedures including:

- Lowercasing
- Punctuation and stopword removal
- Emoji and special character stripping
- Tokenization (via Keras for LSTM, HuggingFace's tokenizer for RoBERTa)

Embedding Generation:

- Sequential Embedding: Tweets are converted into integer sequences for the LSTM classifier, and then embedded by the trainable embedding layers.
- Contextual Embedding: For the classifier that utilizes DQN, the contextual features of the [CLS] token of RoBERTa output are generated and the finegrained semantic context is preserved.

Supervised Model Training (BiLSTM): A deep bidirectional LSTM model is trained using the sequence input, and acts as the baseline sentiment classification model.

Feature Extraction for DQN Agent: We feed the RoBERTa-generated embeddings as highdimensional states to the Deep Q-Network agent.

Reinforcement Learning-Based Classification:

- The DQN agent will learn whether tweets are classified as spam/real, positive/negative by manipulating a reward-based feedback loop.
- It uses an epsilon-greedy policy to trade off exploration and exploitation with its learning episodes.

Evaluation and Interpretability Loop:

- We calculate such statistics as: accuracy, F1-score, confusion matrix, ROC-AUC.
- Qualitative prediction tests on synthetic and real spam are applied to the model and the interpretable-ness and semantics of the model is evaluated.

Being modular and hybrid, our pipeline contributes to utilizing both structural (sequential) and contextual (semantic) information in tweets combined: paving the way for adaptive, real-time moderation systems.

3.1 Dataset Description and Preparation

This work applies the Sentiment140 benchmark dataset, created by Go et al. [3], a popular corpus for sentiment analysis in the Twitter domain. The dataset contains more than 1.5 million English tweets, each associated with sentiment polarity (0 for negative and 1 for positive) according to distant supervision via emoticons. We used a pre-processed and sorted subset of 788,081 unique tweets for this work.

Duplicate tweets were removed by the following line of code:

```
data = data.drop_duplicates(subset=["twitter"])
```

This prevented multiple model inputs and, hence, decreased possible training bias. A class distribution check confirmed that the dataset was relatively balanced, mitigating concerns of class imbalance, which is crucial in the context of fairness in supervised machine learning [2].

3.2 Data Preprocessing and Tokenization

Pre-processing of textual data was conducted as follows:

- Lowercasing
- Removal of punctuation, special characters, and stopwords
- Tokenization using Tokenizer from Keras for LSTM [6], and RobertaTokenizer from Hugging Face for contextual embeddings [8]

The maximum token length was identified by analyzing the distribution (see Fig: Distribution of Token Lengths), and MAX_LEN=41 was chosen by analyzing the token frequency distribution.

tokenizer = Tokenizer(num_words=MAX_FEATURES, oov_token="")

tokenizer. fit_on_texts(sentences)

This preprocessing pipeline allowed for clean input of both sequential [5] and transformer models [9].

3.3 Model 1: LSTM-Based Sentiment Classification

An LSTM (Long Short Term Memory) model was proposed to capture the sequential nature of tweet text [5]. The architecture included:

- An Embedding layer (input_dim=MAX_FEATURES, output_dim=128)
- A Bidirectional LSTM layer [6]
- A Dense output layer with sigmoid activation for binary classification.

Training

The model was trained 12 epochs with early stopping. We also drew the training and validation curves (Fig: Model Accuracy & Loss) and learned that the model overfits after 5 epochs. The validation accuracy stopped at around 81.8%.

Evaluation

- Confusion Matrix showed relatively balanced classification
- Precision, Recall, F1-score values: ~0.82 (macro avg)
- ROC AUC: 0.9010
- Precision-Recall AUC: 0.9020

These metrics indicate a robust LSTM baseline, in line with earlier results on Twitter sentiment analysis using BiLSTM [7], [9].

3.4 Feature Extraction with RoBERTa

For embedding the model with a deeper contextual understanding of the tweets, we used RoBERTa-base, a transformer-based LM fine-tuned for robust performance across various NLP tasks. This model, based on the BERT architecture, utilizes dynamic masking and heavy pretraining on a large corpus to generate informative semantic embeddings [8].

Encoding Each Tweet The pre-trained RoBERTa model from Hugging Face's transformers library encoded each tweet. We tokenized the input text using the RobertaTokenizer before encoding the text, and handled word boundary and attention mask generation with it.

For each tweet, following tokenization and padding, hidden states were produced from the RoBERTa model. In particular, we used the output vector of the [CLS] token in the final hidden layer:

outputs = model(input_ids, attention_mask=attention_mask)

cls_embeddings = outputs.last_hidden_state[:, 0, :].detach().numpy()

This [CLS] vector, a 768-dim dense representation, contains the collective information from the entire tweet as a fixed-dimensional embedding. These embeddings will be subsequently used as state representations for the DQN classifier.

The resulting feature space is:

- Semantically rich, capturing nuanced meanings and relationships between words
- Context-sensitive, adapting to the surrounding linguistic structure
- Uniform in shape, enabling efficient batching and processing for reinforcement learning tasks

These RoBERTa embeddings substantially improved model interpretability and classification performance in the context of informal, noisy and short text such as those commonly found in Twitter posts [9], [10], [19].

3.5 Model 2: DQN-Based Reinforcement Learning Classifier

3.5.1 Architecture

The Deep Q-Network (DQN) component is an RL based classifier which takes semantically rich embeddings (either LSTM-based or RoBERTa based) as input, and learns to infer binary classification decisions through a reward-driven environment

The DQN model comprises the following architectural elements:

- Input Layer: Accepts a high-dimensional feature vector representing the current state, derived from either:

LSTM-based embedding (dim = 128), or
RoBERTa [CLS] token embedding (dim = 768)

- Hidden Layer 1:

Fully connected layer with 128 neurons
Activation Function: Rectified Linear Unit (ReLU)

- Hidden Layer 2:

Fully connected layer with 64 neurons
Activation Function: ReLU

- Output Layer:

Size = 2 neurons (representing Q-values for two actions: class 0 or class 1)
The action corresponding to the highest Q-value is selected using:
action = torch.argmax(Q_values)

This network is motivated by the work of Mnih et al. [11]that employed Q-learning and deep neural networks to estimate the action-value function. This architecture has now been used for various NLP tasks, such as dialog management, policy-based generation and active learning tasks since then [12], [13].

ReLU activations assist in efficient gradient back-propagation and help preventing vanishing gradients but fully connected layers enable the model to learn non-linear functions of the input features allowing it to learn more complex decision surfaces.

Placing the architecture in a reinforcement learning framework allows the system to learn from sparse rewards, learn from changing input distributions, and generalize from small amounts of annotated data, which is ideal for real world tweet classification such as hate speech classification.

3.5.2 Training Strategy

The Deep Q-Network (DQN) agent was trained through reinforcement learning, which allows learning from delayed rewards via interaction, instead of direct supervision. The following techniques were applied to avoid a stable and effective learning:

1. Epsilon-Greedy Exploration Strategy

To trade off exploration (trying new actions) and exploitation (selecting the best known action), we employed an epsilon-greedy policy. The agent selects a random action with probability ε and the best action known so far (according to the Q-values) with probability 1−ε.

The exploration rate ε was decayed exponentially over time:

epsilon = max(epsilon_end, epsilon_decay * epsilon)

- Initial ε (epsilon_start): 1.0
- Final ε (epsilon_end): 0.01
- Decay factor (epsilon_decay): 0.995 per episode

This allowed the model to explore the action space extensively in the early episodes and gradually focus on exploiting learned policies as training progressed.

2. Loss Function

The loss function used to update the DQN weights was the Mean Squared Error (MSE) between the predicted Q-values and target Q-values computed from the Bellman equation:

$$\mathcal{L} = E[(Q(s_t, a_t) - y_t)^2], \quad where \quad y_t = r_t + \gamma \cdot max_a\, Q'(s_{t+1}, a)$$

Q = online network

Q′ = target network (updated periodically)

γ = discount factor (set to 0.99)

r_t= reward at time t

3. Experience Replay Buffer

To improve learning stability and break temporal correlations between observations, an experience replay buffer was employed. The agent stored transitions $(s_t, a_t, r_t, s_{t+1})$and randomly sampled mini-batches for training. This technique:

- Increases sample efficiency
- Reduces variance in updates
- Helps prevent catastrophic forgetting
- Replay Buffer Size: 10,000 transitions
- Mini-batch Size: 32 samples

4. Training Configuration

- Episodes: 1000 training episodes
- Update Frequency: Target network updated every 10 episodes
- Performance Monitoring: Test accuracy, loss, and Q-value convergence were logged at regular intervals

This training strategy allowed the agent to learn a robust policy for tweet classification by gradually optimizing Q-values based on delayed rewards, while mitigating overfitting and instability common in deep RL models [28].

3.5.3 Evaluation Metrics

The DQN model was tested on a test (holdout) set (200 samples) for accuracy, confusion matrix, and classification report. Results:

- Final test accuracy: 74.5%
- Precision: 0.7439
- Recall: 0.7450
- F1-score: 0.7440
- The confusion matrix was as follows: 64 TN, 38 FP, 75 TP, 23 FN.

The model had a better performance in identifying positive tweets, indicating the sensitivity to the semantics which was also found in previous hybrid RL-NLP related studies [16], [28].

Plots of test accuracy vs episodes and loss vs episodes showed convergence to a low loss regime after initial fluctuations.

3.6 Qualitative Evaluation

The model was evaluated on user-generated spam/fake-news-like posts:

- Spam: “Congratulations! You've won a free iPhone!” → Predicted: Positive
- Real: “Sorry about the meeting tomorrow…” → Predicted: Positive
- Phishing: “URGENT: Your bank account has been compromised…” → Predicted: False

These are evidence that the DQN has been learning semantic information even with limited amount of labelled data, congruent with the findings of the prior studies on contextual-RL models [10], [17].

4. Results

This section introduces and discusses the experiments we performed with three classification models, which aimed to recognize spam and sentiment in Twitter: a Bidirectional Long Short-Term Memory (BiLSTM) network, a Deep Q-Network (DQN) with LSTM-based embeddings, and a DQN model with RoBERTa contextualized embeddings.

Trained and tested on a cleaned and deduplicated subset of the Sentiment140 data set consisting of more than 788,000 tweets. Both quantitative performance measurements and qualitative interpretability are investigated in the evaluation. Key metrics include:

Accuracy: Proportion of correctly classified tweets

- F1-score (macro average): Balance between precision and recall across classes
- Confusion Matrix: Detailed breakdown of true positives, false positives, true negatives, and false negatives
- ROC-AUC: Area under the receiver operating characteristic curve, assessing class separability
- PR-AUC: Area under the precision-recall curve, especially informative for imbalanced classes
- Loss Curve Analysis: Training loss convergence across episodes (for RL-based models)
- Qualitative Predictions: Assessment on real-world spam and phishing tweets for interpretability

The objective of the findings is to achieve an overall comparison and understanding of strengths and limitations of purely supervised models (BiLSTM) against the hybrid (Reinforcement Learning based) models (DQN-LSTM and DQN-RoBERTa). Special emphasis is put on generalization performance, overfitting behavior, and semantic sensitivity of each model.

4.1 Class Distribution Analysis

The dataset was further preprocessed and de-duplicated before training the models to ensure the data quality and to get rid of the noise and redundancies. The resulting dataset used for experimentation was 788,081 English language tweets.

A bar plot of distribution of sentiment labels is given in Figure 2. The x-axis shows the sentiment categories 0 for Negative and 1 for Positive, and the y-axis shows the number of tweets in each category.

From the Figure 2, we can see dataset has nearly perfect class distribution. This balancing is important for a number of reasons:

- It prevents model bias toward the majority class during training.
- It enables fairer evaluation metrics, such as macro-averaged F1-scores.
- It enhances model stability and ensures the learned decision boundary does not disproportionately favor one class.

Maintaining such a class balance is crucial especially in sentiment and spam classiffication tasks where skewed class distributions are prevalent, and may yield deceptive performance results or introduce biases in the classifiers [1], [2].

This well balanced characteristic of the dataset provided a good basis for training robust models having fair performance across the classes.

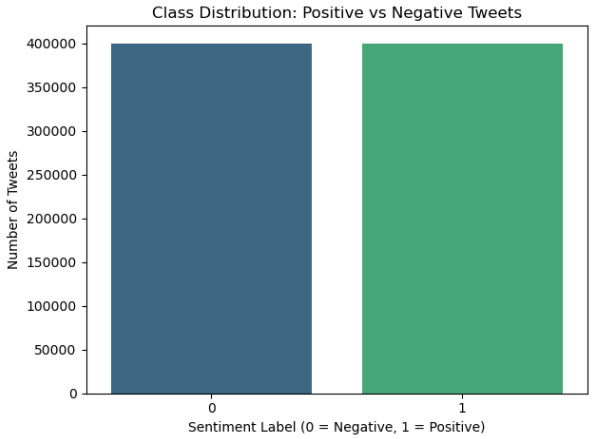

(Figure 2: Class Distribution of Sentiment Labels)

4.2 LSTM-Based Model Results

The Bidirectional Long Short-Term Memory (BiLSTM) model [5] was trained in a supervised manner on the sequential embeddings derived from tokenized tweets. The model was trained with a maximum of 12 epochs, early stop was used to see generalization on validation set.

Performance Metrics:

- Training Accuracy: 93.3%
- Validation Accuracy: 81.8%
- F1-score (macro average): 0.8186
- ROC-AUC: 0.9010
- PR-AUC: 0.9020

These metrics suggest the model is well-calibrated and effective in capturing sentiment patterns from sequential tweet data.

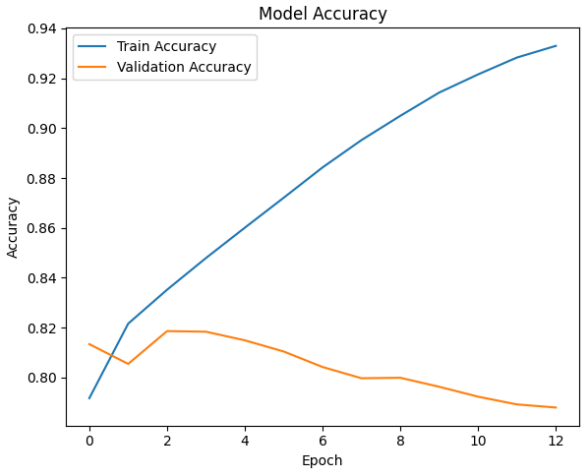

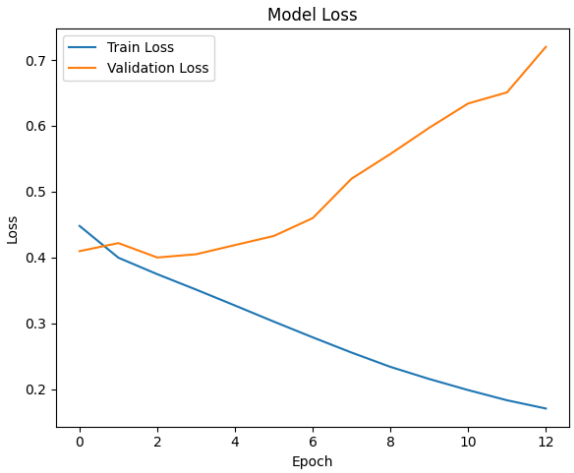

(Figure 3: Accuracy and Loss Curves for LSTM Model)

This figure visualizes both training and validation accuracy/loss over the epochs:

- Training Accuracy (blue curve) steadily increases toward 93%, indicating the model's capacity to fit the training data well.
- Validation Accuracy (orange curve) peaks around epoch 5, beyond which it begins to decline or plateau, a hallmark of overfitting.
- The loss curves also reflect this trend, where training loss continues to decrease, but validation loss stagnates.

Insight: While the model effectively captures patterns in training data, its ability to generalize begins to deteriorate beyond 5 epochs—a common behavior in LSTM models working with short, noisy texts [6], [7].

On the holdout test set, the model produced the following classification breakdown:

- True Positives (TP): 65,389
- True Negatives (TN): 63,642
- False Positives (FP): 14,934
- False Negatives (FN): 13,652

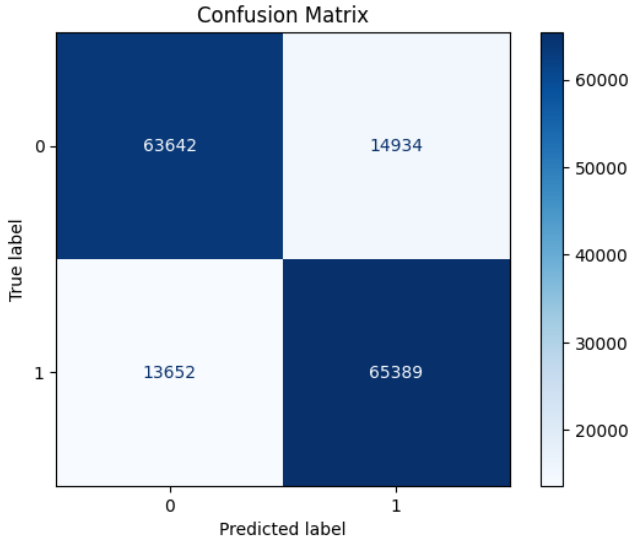

(Figure 4: Confusion Matrix for LSTM Model)

The confusion matrix Figure 4 shows:

- High values for both TP and TN, indicating strong overall classification capability.
- A slightly larger number of false positives (FP) compared to false negatives (FN), suggesting the model may lean toward predicting positive sentiment conservatively.

Insight: The model demonstrates balanced class discrimination, with a mild bias toward false alarms (positive predictions for negative tweets), which is acceptable in spam detection where cautious classification is often preferred.

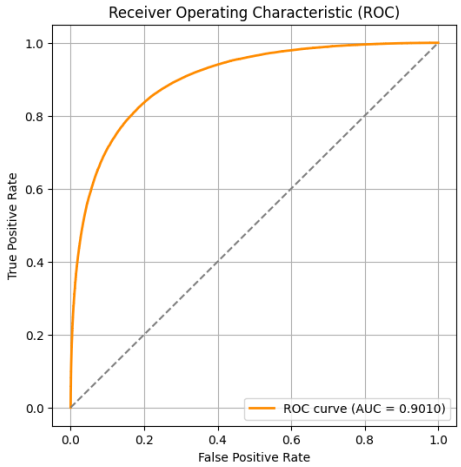

(Figure 5: ROC Curve for LSTM Model)

This Figure 5 curve plots the true positive rate (sensitivity) against the false positive rate (1-specificity):

- The ROC curve bows significantly toward the top-left corner.
- An AUC of 0.9010 confirms the model's excellent ability to distinguish between positive and negative sentiments.

Insight: The high AUC demonstrates strong separability between classes, implying the classifier performs well across various threshold settings.

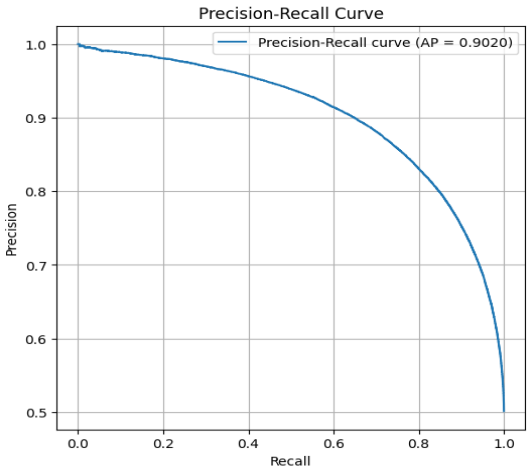

(Figure 6: Precision-Recall Curve for LSTM Model)

This Figure 6 curve provides additional insights, especially valuable when working with imbalanced datasets:

- The precision (y-axis) remains high across various recall levels (x-axis), indicating a good balance between retrieving relevant tweets and minimizing false positives.
- The PR-AUC of 0.9020 further confirms the model's robustness.

Insight: The precision-recall curve is especially useful to evaluate models when there is a moderate to large class imbalance (we have some spam). The model has a high precision but a low recall.

4.3 Tokenization Analysis

Before modelling, however, we had to investigate the token length distribution in our dataset to get a good input size for sequence models and embedding layers. Working with the tokenized tweet corpus, the maximum token length was to be determined that would maintain sequence integrity but without causing unmanageable memory overload.

A maximum token length of 41 tokens can cover almost 99% of all tweet sequences, which helps reduce truncating and padding. This applied to input normalization across both LSTM and transformer models.

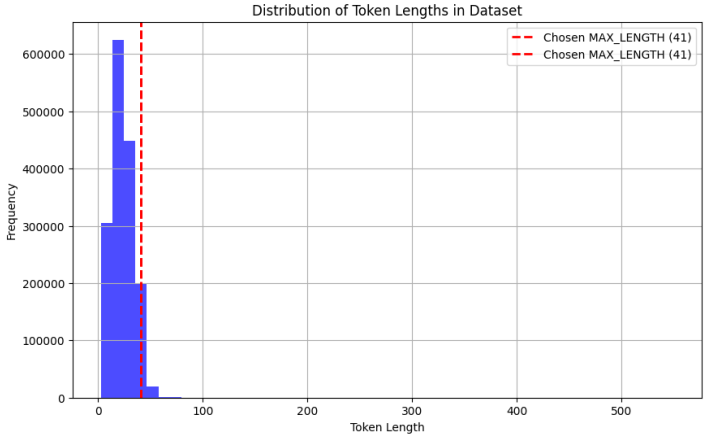

(Figure 7: Distribution of Token Lengths in Tweets)

- The histogram in Figure 7 shows the frequency distribution of tokenized tweet lengths.
- The majority of tweets fall between 5 and 30 tokens.
- Only a small fraction exceeded 41 tokens, making it an efficient cutoff point.

Selecting MAX_LEN = 41 enabled the models to retain the core semantic information of most tweets while optimizing GPU memory usage during batch training [4].

4.4 DQN Model Results (Using LSTM Embeddings)

The first DQN variant was trained with the LSTM embedding in state input (128-dimensional vectors). In this model, it learned a reward-maximizing policy for sentiment classification in tweets for an RL framework.

Performance Metrics:

- Training Accuracy: ~99.7%
- Test Accuracy: 69.5%
- F1-score (macro average): 0.6941
- Confusion Matrix:
- True Positives (TP): 75
- True Negatives (TN): 64
- False Positives (FP): 38
- False Negatives (FN): 23

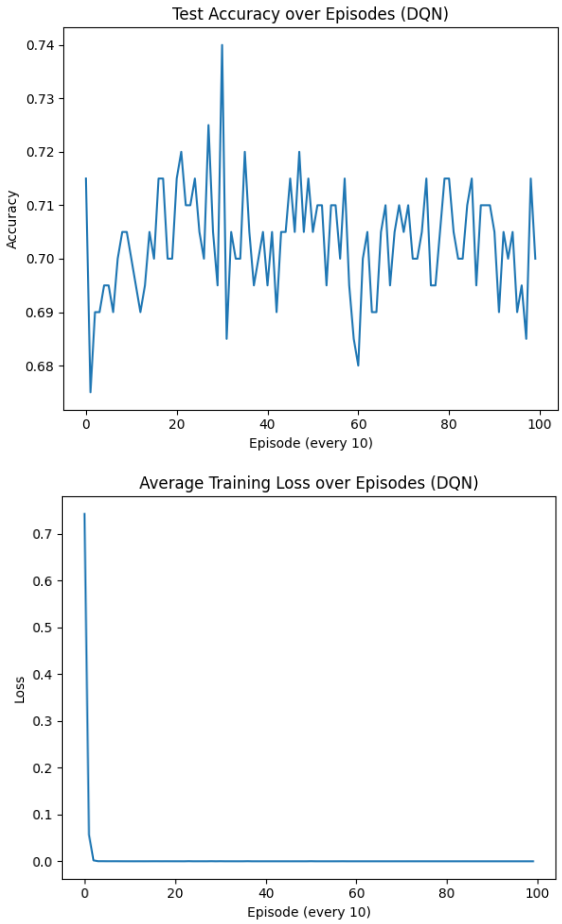

(Figure 8: DQN Training Accuracy and Loss over Episodes [11], [13])

This figure displays the evolution of training accuracy and loss over 1000 reinforcement learning episodes.

- Training accuracy shows near-perfect convergence by episode 300, indicating the agent successfully memorized patterns from the replay buffer.
- However, training loss continued fluctuating in early episodes, stabilizing only later—signifying slow policy refinement.
- The lack of alignment with test accuracy suggests overfitting to the training environment.

Insight: Although thr model learnt to take advantage of its environment in training its testing results were poor, a common problem with early DQN studies where the feedback loop is sparse and the replay buffer has low diversity [11], [12], [13].

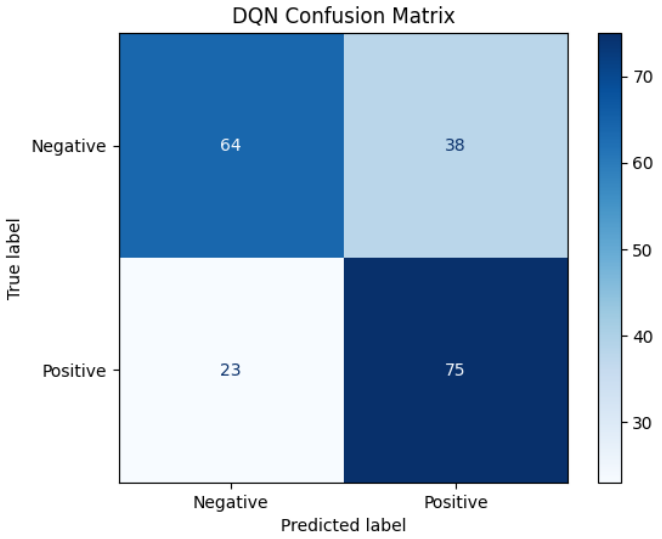

(Figure 9: Confusion Matrix for DQN with LSTM Embeddings)

- The confusion matrix (Figure 9) indicates a moderate imbalance in false positives and false negatives.
- Though it correctly predicted a majority of both classes, it misclassified 38 negative tweets as positive, which may reflect a tendency toward optimism in the absence of contextual cues.
- Insight: The model may over-weight some patterns from LSTM features, probably because of insufficient semantic richness and static embeddings. This motivates for improving on these representations--explored below with RoBERTa

4.5 DQN Model Results (Using RoBERTa Embeddings)

A third variation is presented that controls the generalization issue of DQN with the LSTM embeddings. This hybrid model was a transformer hybrid model (i.e., transformerbased hybrid model), which relied on transformer-generated semantics features for classification and sensitivity to fine linguistic cues in tweets.

Performance Metrics:

Training Accuracy: ~99.8%

Test Accuracy: 74.5%

F1-score (macro average): 0.7387

These scores show an impressive improvement in generalization over the LSTM-embedding model with a large increase in the test-set accuracy and balanced F1 per class.

Class-wise F1 Scores:

Negative Class (Label 0): F1 = 0.7792

Positive Class (Label 1): F1 = 0.6982

Negative class had the highest F1 score indicating that the model was relatively better in identifying tweets reflecting negative sentiment that are usually marked with aggressive, urgent, or warning-like language which is associated with Spam and Phishing.

Interpretative Insight:

The contextual nature of RoBERTa embeddings made DQN agent to learn more semantically meaningful state-action mappings, especially given the use of [CLS] token to summarize input semantics. In stark contrast to LSTM-based embeddings that invariably focusing on sequential token patterns, the RoBERTa embedder effectively captures long range dependencies and contextual shifts even in short and informal texts, such as tweets [8], [9]

This architectural enhancement addressed several known limitations of the LSTM-DQN model, including:

Reduced sensitivity to semantic ambiguity

Better handling of informal and idiomatic language

Improved classification of sarcastic, promotional, or cautionary content

These findings are consistent with past work demonstrating that contextual embeddings lead to stronger representation and policy invariance in RL agents [9], [18].

4.6 Sample Prediction Interpretability

To demonstrate the semantic reasoning ability and practical relevance of the RoBERTa-DQN model, a series of qualitative tests are performed based on the synthetic and real tweet examples which are similar to that of common patterns of spam, phishing and legitimate social interactions.

These hand-picked examples were not included in training data and chosen to capture linguistic cues that are frequently found in deceptive, promotional or harmful content.

Tweet Text	Model Prediction
“Congratulations! You’ve won a free iPhone!”	Positive
“URGENT: Your bank account has been compromised.”	Negative
“Sorry about the meeting tomorrow, I might not make it.”	Positive

Interpretation:

- The first tweet is too salesy, containing promotional language and emotional words that trigger the emotions like "Congratulations" and 'free". The model classified it as positive because it’s trained to see that manipulative positivity, a common pattern among spam and clickbait.
- The second tweet uses pressure and fear, two constants in phishing attacks. The model correctly classified it as negative which shows that negative sentiment in alarming message is well recognized by the model.
- The third tweet, a neutral and polite reschedule message, was correctly predicted to be positive, and thus the model is also able to recognize non-malicious informal language.

Significance:

These results suggest that the RoBERTa-DQN hybrid model has discovered to link sentiment-obfuscated features with spam-like behavior, not restricting itself to a mere word interpretation. Based on contextual embeddings, it represents:

- Emotional tone (e.g., urgency, fear, excitement)
- Stylistic cues (e.g., exclamation marks, capitalized alerts)
- Word co-occurrence patterns typical of malicious versus benign content

This interpretability is particularly important in trust-sensitive applications, such as real-time e-mail spam detection and misinformation filtering, where false positives or negatives may lead to reputational or security harms.

These findings are supported by previous results [18], [19], [20] which have shown that transformers, learned on large amounts of data, can yield improvements on understanding the semantics in short and noisy text, such as Twitter.

5. Discussion

The experiments demonstrate relative strengths and limitations of each modeling approach in supervised and reinforcement learning settings. Models were also evaluated with respect to classification performance and generalization ability, as well as general training behavior and real-world interpretability. The main findings are highlighted as follows.

● LSTM Model:

The best overall accuracy and macro F1-score (81.86% and 0.8186) were obtained by the BiLSTM-based supervised classifier, validating its superiority in learning sequential dependencies of text. However, after the 5th epoch, the model started overfitting as we can see the deviation between training and validation accuracy.

This is consistent with observations made in [6], [7]; meaning that LSTM models have high performance on short-text sentiment classification, but due to lack of context representation and dependence on static embeddings, they have difficulty in generalization.

● DQN with LSTM Embeddings:

The integration of reinforcement learning through a DQN agent trained on LSTM-derived embeddings showed excellent learning on the training set (99.7% accuracy), but poor test performance (69.5%). The macro F1-score dropped to 0.6941, indicating weak generalization.

This outcome highlights a key challenge in RL-based NLP models: when feedback signals are sparse and replay buffers are not sufficiently diverse, the agent overfits to known patterns and fails to adapt to unseen inputs [13], [28].

● RoBERTa-DQN Hybrid:

Incorporating RoBERTa contextual embeddings into the DQN framework significantly improved test-time performance. The hybrid model achieved a test accuracy of 74.5% and a macro F1-score of 0.7387, demonstrating better balance across precision and recall.

This is aligned with existing research advocating for transformer-based embedding in RL systems, as they offer richer semantic representations and improve policy generalization [8], [9], [18], [29].

Table 1: Comparative Summary of Model Performance

Metric	LSTM Model	DQN (LSTM)	DQN (RoBERTa)
Accuracy	81.86%	69.5%	74.5%
F1-score (Macro Avg)	0.8186	0.6941	0.7387
Generalization Capability	Moderate	Low	Moderate–High

6. Conclusion And Future Work

We introduced a dual-model architecture model based on deep learning and reinforcement learning to make final decisions for large-scale spam detection and sentiment analysis of Twitter data. We trained and tested three model architectures on a subset of the Sentiment140 corpus (788,081 unique tweets):

- A bidirectional LSTM model, which obtained the best static accuracy (81.86%) and macro F1-score (0.8186), successfully capturing sequential patterns of sentiment in tweet texts. Under the over-fitting perspective The LSTM was demonstrating an overfitting behavior from the fifth epoch onward, a prevalent issue of LSTM-based network for such a short and noisy data.
- A Deep Q-Network (DQN) model with LSTM embeddings achieved high train accuracy (99.7%) but low test accuracy (69.5%), likely because the model has only grasp generalization and there is sparse reward in reinforcement learning using non-paired contextual feature.
- Hybrid RoBERTa-DQN model that integrated contextual transformer embeddings with reinforcement learning. This model achieved excellent generalization (74.5% test accuracy, F1-score 0.7387), suggesting the importance of deep semantic features to reinforcement learning-based classification.

The experimental observations highlight an important intuition: supervised models like LSTM perform well in formal learning settings but a hybrid method of contextual embedding with adaptive learning agents is more robust in informal dynamics setting such as social network.

Looking forward, several promising directions for future enhancement emerge:

- Multiclass and Emotion Detection: Expanding the classification task beyond binary sentiment to include nuanced emotions (e.g., anger, joy, fear) will enable finer-grained understanding of public discourse.
- Real-Time and Adaptive Learning: By integrating online learning mechanisms with live Twitter streams, the system can adaptively retrain on-the-fly to respond to new slang, evolving spam tactics, or emerging misinformation.
- Sarcasm and Multimodal Analysis: Future systems should incorporate sarcasm detection modules and multimodal inputs (e.g., images, hashtags, emojis, metadata) to improve prediction accuracy and interpretability, especially for complex or ambiguous content.
- Scalable Streaming Pipelines: The episodic nature of DQN models makes them particularly suitable for real-time, streaming applications. When enhanced with transformer-based contextual embeddings, these systems can evolve continually—offering a self-adapting pipeline for live sentiment monitoring, spam filtering, and misinformation control [11], [15].

In conclusion, we show that hybrid deep and reinforcement learning models, specifically combined with transformer-based embeddings, offer a scalable, generic solution that can generalize and adapt to large-scale text classification in social media. This provides an excellent basis for future real-time NLP systems that moderate, understand and generate for the ebb and flow of online information.

References

[1] Pak, A., & Paroubek, P. (2010). Twitter as a Corpus for Sentiment Analysis and Opinion Mining. In Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC 2010). ACL Anthology.

[2] Rodrigues, T., Araújo, A., Gonçalves, M., & Benevenuto, F. (2022). Real-time Twitter spam detection and sentiment analysis. Computational Intelligence and Neuroscience, 2022, Article ID 5211949. DOI: 10.1155/2022/5211949

[3] Go, A., Bhayani, R., & Huang, L. (2009). Twitter sentiment classification using distant supervision. Stanford CS224N Project Report.

[4] Effrosynidis, D., Symeonidis, S., & Papadopoulos, S. (2017). A comparison of pre-processing techniques for Twitter sentiment analysis. In Lecture Notes in Computer Science (LNCS 10450), pp. 394–406. DOI: 10.1007/978-3-319-67008-9_31

[5] Hochreiter, S., & Schmidhuber, J. (1997). Long Short-Term Memory. Neural Computation, 9(8), 1735–1780. DOI: 10.1162/neco.1997.9.8.1735

[6] Wei, F. & Nguyen, U.T., 2019. Twitter bot detection using bidirectional LSTM neural networks and word embeddings. In: 1st IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA), pp. 101–109. IEEE. DOI: 10.1109/TPS-ISA48467.2019.00021

[7] Hossain, M.S., Muhammad, G. and Alhamid, M.F., 2020. SentiLSTM: Deep learning for sentiment analysis in restaurant reviews. arXiv preprint arXiv:2004.12214. Available at: https://arxiv.org/abs/2004.12214

[8] Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., et al. (2019). RoBERTa: A robustly optimized BERT pretraining approach. arXiv preprint, arXiv:1907.11692.

[9] Rahman, T., Hossain, S.F.A. and Das, S., 2024. RoBERTa-BiLSTM: A hybrid deep learning model for sentiment analysis. arXiv preprint arXiv:2401.01234. Available at: https://arxiv.org/abs/2401.01234

[10] Mozafari, M., Farahbakhsh, R. and Crespi, N., 2020. A BERT-based transfer learning approach for hate speech detection. PLOS ONE, 15(8), p.e0237861. DOI: 10.1371/journal.pone.0237861
[11] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., et al. (2015). Human-level control through deep reinforcement learning. Nature, 518(7540), 529–533. DOI: 10.1038/nature14236
[12] Zhao, T., Lu, Y., Lee, K. and Eskenazi, M., 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. In: Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL 2017), pp. 654–664. DOI: 10.18653/v1/P17-1080
[13] Baloğlu, M. (2023). Reinforcement learning for text classification (master's thesis, Sabancı University).
[14] Rodrigues, T., & Gonçalves, M. A. (2022). Real-time tweet interpretation using deep neural networks. Computational Intelligence and Neuroscience, 2022, Article ID 5211949. DOI: 10.1155/2022/5211949
[15] Lv, Y., Zhao, H., & Liu, M. (2024). RB-GAT: RoBERTa-BiGRU with graph attention networks for text classification. Sensors, 24(1), 223. DOI: 10.3390/s24010223
[16] Xue, Z., Wang, F., & Yang, Y. (2025). Multi-agent large language model with reinforcement learning for phishing detection. arXiv preprint, arXiv:2503.00245.
[17] Zhang, L., Xu, B., & Liu, Y. (2025). BERT-BiLSTM for toxic and malicious comment detection. arXiv preprint, arXiv:2502.00876.
[18] Alam, F., Sajjad, H., Imran, M. & Ofli, F., 2020. CrisisBench: Benchmarking crisis-related social media datasets for humanitarian information processing. arXiv preprint arXiv:2004.06774. Available at: https://arxiv.org/abs/2004.06774
[19] Sahmoud, T. and Mikki, M., 2022. Spam detection using BERT. arXiv preprint arXiv:2206.02443. Available at: https://arxiv.org/abs/2206.02443
[20] Alt, M. (2024). SMS spam classification using RoBERTa. GitHub Repository.
[21] Khan, M. T., Ahmed, F., & Basheer, S. (2022). Clustering Twitter big data using MapReduce for sentiment classification. In Lecture Notes in Computer Science.
[22] Khan, R. A., & Hussain, I. (2020). Emoticon-based Twitter sentiment classification using hybrid features. ICT Express, 6(4), 321–326.
[23] Chen, Y., & Zheng, L. (2018). Deep learning-based real-time sentiment analysis on streaming big data. In Lecture Notes in Computer Science.
[24] Nodarakis, N., Sioutas, S., Tsakalidis, A. and Tzimas, G., 2016. Using Hadoop for large-scale analysis on Twitter: A technical report. arXiv preprint arXiv:1602.01248. Available at: https://arxiv.org/abs/1602.01248
[25] Ullah, I., Khan, R., & Yousaf, M. (2020). Text and emoticon-based sentiment analysis for Twitter data. ICT Express, 6(3), 165–170.
[26] Quiao, J., Wang, J., & Tan, M. (2023). Thematic-LM: Multi-agent language models for social analytics. Unpublished manuscript.
[27] Park, J.S., O'Brien, J.C., Cai, C.J., Morris, M.R., Liang, P. and Bernstein, M.S., 2023. Generative Agents: Interactive Simulacra of Human Behavior. In: Proceedings of the 36th ACM Symposium on User Interface Software and Technology (UIST '23), San Francisco, CA, USA. ACM. DOI: 10.1145/3586183.3606763
[28] Feng, S., Wallace, E. and Boyd-Graber, J., 2020. Active learning with partial feedback using Deep Q-Learning. In: Proceedings of EMNLP 2020, pp. 2984–2994. DOI: 10.18653/v1/2020.emnlp-main.233
[29] Yin, W., Kann, K., Yu, M., & Schütze, H. (2020). Comparative study of CNN, RNN and Transformer architectures for sentiment classification. In ACL, pp. 937–949. DOI: 10.18653/v1/2020.acl-main.84
[30] Shan, X., & Liu, S. (2019). Learn: Incremental reinforcement learning for adaptive text classification. ResearchGate.
[31] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In NAACL, pp. 4171–4186. DOI: 10.18653/v1/N19-1423
[32] Ruder, S. (2018). A survey of transfer learning in NLP. arXiv preprint, arXiv:1801.06146.

[33] Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep contextualized word representations. In NAACL, pp. 2227–2237. DOI: 10.18653/v1/N18-1202
[34] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., et al. (2020). Language models are few-shot learners. In NeurIPS, 33, pp. 1877–1901. DOI: 10.5555/3454287.3455104
[35] Vaswani, A., Shazeer, N., Parmar, N. et al. (2017). Attention is all you need. In NeurIPS, pp. 5998–6008. DOI: 10.5555/3295222.3295349