

UNIVERSITY OF BAHRAIN

College of Engineering



**Department of
Electrical & Electronics
Engineering**

ARTIFICIAL INTELLIGENCE BASED SENTIMENT ANALYSIS FOR SOCIAL MEDIA

A Thesis Submitted in Partial Fulfilment of the Requirements for the
M.Sc. Degree in Artificial Intelligence Systems

Submitted by
Ammar Yaser Maki Isa
ID: 20174295

Supervised by
Dr. Ayman Al-Khazraji
Assistant Professor
University of Bahrain

Kingdom of Bahrain
November 2025

UNIVERSITY OF BAHRAIN

College of Engineering



Department of
Electrical & Electronics
Engineering

Defence committee's approval

The student Ammar Yaser Maki Isa, ID 20174295, has defended the thesis by which he fulfilled the requirement for the degree of Master of Science entitled "Ai-Based Sentiment Analysis in Social Media".

Name:

Dr. Ayman Al-Khazraji (Supervisor)

A blue ink signature of Dr. Ayman Al-Khazraji, written in a cursive style.

Prof. Salam Ismaeel (External Examiner)

A blue ink signature of Prof. Salam Ismaeel, written in a cursive style.

Dr. Mohammed Alkhalidy (Internal Examiner)

A blue ink signature of Dr. Mohammed Alkhalidy, written in a cursive style.

Academic Department Stamp



Abstract

The exponential growth of social media has established platforms accessed through public APIs as key venues for public discourse and sentiment expression. This research presents a comprehensive AI-based system for real-time sentiment analysis of social-media streams using advanced transformer-based language models. To address challenges related to noisy, short, and multilingual the system integrates a robust preprocessing pipeline capable of handling linguistic variability, and informal writing patterns. The primary objective is to develop an end-to-end sentiment-classification platform that not only performs real-time inference but also enables users to experiment with, evaluate, and fine-tune multiple sentiment-analysis models through an intuitive graphical interface. Rather than focusing on isolated datasets, the methodology leverages diverse large-scale social-media sources and multilingual input streams to ensure realistic, generalisable performance. The system architecture incorporates multiple model categories: traditional deep-learning models such as LSTM and CNN for lightweight inference, and transformer-based models including RoBERTa and XLM-RoBERTa for enhanced multilingual classification. A unified model management layer enables seamless switching between locally trained models and pretrained transformer models across the application. The resulting platform supports real-time sentiment monitoring, multilingual analysis, model training and evaluation, transformer fine-tuning, and interactive visualisations, providing an accessible research environment suitable for users.

Acknowledgements

I am deeply grateful to my supervisor, Dr. Ayman Al Khazraji, for his guidance and mentorship, which have been instrumental in the completion of this master's thesis. I would like also to dedicate this thesis to my loved ones, including my family, friends, and colleagues, who have supported me throughout this journey and filled my heart with love and motivation.

Table of Contents

Defense Committee’s approval.....	A
Abstract.....	C
Acknowledgements.....	D
Table of Contents.....	E
List of Tables	H
List of Figures	I
List of Abbreviations	K

Chapter One

Introduction to the Study

1.1 Introduction.....	1
1.2 Research Problem	3
1.3 Research Questions	4
1.4 Research Objectives.....	5
1.5 Research Significance.....	6
1.6 Study Setting.....	6
1.7 Research Terminologies.....	7

Chapter Two

Literature Review

2.1 Overview of Sentiment Analysis	9
2.1.1 Historical Development of Sentiment Analysis.....	9
2.1.2 Granularity Levels and Applications of Sentiment Analysis.....	10
2.1.3 Challenges in Social Media Sentiment Analysis	12
2.2 Traditional Approaches to Sentiment Analysis	13
2.2.1 Lexicon-Based Methods	13
2.2.2 Machine Learning Approaches	14
2.2.3 Feature Engineering for Sentiment Analysis	15
2.2.4 Recurrent Neural Networks (LSTM/GRU)	16
2.2.5 Convolutional Neural Networks for Text	18
2.2.6 Word Embeddings and Contextual Representations.....	19
2.3 Transformer Architectures	20
2.3.1 Attention Mechanism and Self-Attention	20
2.3.2 BERT and Its Architecture.....	23

2.3.3 RoBERTa and Optimization Techniques.....	25
2.3.4 XLM-RoBERTa Overview	27
2.4 Fine-Tuning Strategies for Sentiment Analysis.....	30
2.4.1 Evolution of Transfer Learning Approaches	30
2.4.2 Domain Adaptation Techniques	31
2.4.3 Few-Shot and Zero-Shot Learning.....	33
2.5 Multilingual Sentiment Analysis	35
2.5.1 Cross-Lingual Transfer Learning.....	35
2.5.2 Language-Specific Challenges.....	37
2.5.3 Recent Advances in Multilingual Models.....	39
2.6 Performance Evaluation.....	40
2.6.1 Metrics for Sentiment Analysis	40
2.6.2 Benchmark Datasets.....	42
2.6.3 Recent Comparative Studies	44
2.7 Research Gaps.....	45
2.7.1 Current Limitations.....	45
2.7.2 Potential Applications.....	47

Chapter Three

Methods and Procedures

3.1 System-Architecture Overview.....	49
3.2 Data-Collection Pipeline.....	52
3.3 Datasets	54
3.4 Pre-processing & Augmentation.....	56
3.5 Model Architecture	58
3.5.1 Transformer-Based Transfer Learning	60
3.6 Training and Fine-Tuning Procedure.....	61
3.7 GUI Implementation	64
3.7.1 Page Layout and Navigation.....	64
3.7.2 Page 1 - Single-Tweet Sentiment.....	65
3.7.3 Page 2 - Real-Time Hashtag Monitor	66
3.7.4 Page 3 - Instant Sentiment Analysis	67
3.7.5 Page 4 - Multilingual Stream Monitor	70
3.7.6 Page 5 & 6 - Model Training & Evaluation Dashboard.....	74

3.7.7 Page 7 and 8 settings & fine tuning	79
3.7.8 Threading and Responsiveness	80
3.8 Experimental Setup	81
3.8.1 Version Pinning	81
3.8.2 Evaluation Protocol.....	82
3.9 Evaluation Metrics & Statistical Tests.....	82
3.10 Chapter Conclusion.....	85

Chapter Four

Results, Discussion and Recommendations

4.1 Results of Self-Trained Models	87
4.2 Results of Transfer-Learning Models	92
4.3 Comparative Analysis	94
4.4 Discussion	96
4.5 Recommendation and Future Work	99
4.6 Chapter Summary	100
References.....	102
Abstract (Arabic)	114

List of Tables

Table (2.1): CNN comparisons	19
Table (3.1): Sample of Extracted twitter data	53
Table (3.2): Kaggle Corpus Class Distribution.....	55
Table (4.1): Test-set Metrics for Transfer-Learning Models.....	92
Table (4.2): Per-Language F1 Scores on Live Streams	92
Table (4.3): Metrics Across Self Trained and Transfer Learning Models.....	95
Table (4.4): Research Questions to Key Findings and Evidence.....	98
Table (4.5): Recommendation Matrix.....	100

List of Figures

Figure (1.1): Overview of AI-Based Sentiment Analysis Pipeline	3
Figure (1.2): Research Questions on Sentiment Analysis System.....	5
Figure (2.1): Historical Development of Sentiment Analysis	9
Figure (2.2):Granularity levels of sentiment analysis	11
Figure (2.3): BERT-large Architecture	25
Figure (3.1): System Architecture Diagram.....	50
Figure (3.2) Sentiment Analysis System Architecture	51
Figure (3.3): Data-Collection and Processing Pipeline.....	52
Figure (3.4): Live Stream Statistics by Language	55
Figure (3.5): Pre-processing & Sequence Workflow Diagram.....	57
Figure (3.6): Model Integration in System Architecture	58
Figure (3.7): Neural Network Architecture for Text Classification	59
Figure (3.8): Custom Sentiment Classifier Architecture	59
Figure (3.9): Models Comparison.....	60
Figure (3.10): RoBERTa Transfer-Learning Architecture	60
Figure (3.11): RoBERTa Fine Tuning Performance Curve.....	62
Figure (3.12): Home page and Navigation	64
Figure (3.13): Navigation Menu	65
Figure (3.14): Page 1 Single-Tweet Analysis	66
Figure (3.15): Page 1 Single-Tweet Analysis Drop list	66
Figure (3.16): Page 2 Real-Time Hashtag Monitor	67
Figure (3.17): Page 3 Positive Analysis through WhatsApp	68
Figure (3.18): Page 3 Positive Analysis Results	68
Figure (3.19): Page 3 Negative Analysis through Whatsapp.....	68
Figure (3.20): Page 3 Negative Analysis Results	69
Figure (3.21): Page 3 Settings and Adjustments.....	69
Figure (3.22) Logs Folder contains instant sentiment analysis results	70
Figure (3.23) Page 4 Multilingual Twitter Opinion.....	71
Figure (3.24): Page 4 Sentiment Analysis Visualization Results	71
Figure (3.25): Page 4 Bar Chart Visualization Results.....	72
Figure (3.26): Page 4 Time Trends Visualization Results.....	72
Figure (3.27): Page 4 Language Analysis Visualization Results.....	73

Figure (3.28): Page 4 Time Trends Area Chart Visualization Results	73
Figure (3.29): Page 4 Correlations Visualization Results.....	74
Figure (3.30): Page 4 Sample tweets display	74
Figure (3.31): Page 5 Data Preparation.....	75
Figure (3.32): Page 5 Model Selection	76
Figure (3.33): Page 5 Customize LSTM Model Architecture	76
Figure (3.34): Page 5 Customize CNN Model Architecture.....	76
Figure (3.35): Page 5 Model Training Hyperparameters	77
Figure (3.36): Page 5 Model Evaluation & Comparison	77
Figure (3.37): Page 6 Simplified Model Comparison	78
Figure (3.38): Page 6 Simplified page results	78
Figure (3.39) Text Preprocessing options	79
Figure (3.40): Model change choices.....	80
Figure (3.41): Fine tuning page	80
Figure (3.42). Precision, Recall, F1 Comparison.....	83
Figure (3.43): Confusion Matrix for Model Comparison	84
Figure (4.1): Model Accuracy by Dataset.....	88
Figure (4.2): F1 Score by Dataset	88
Figure (4.3): Inference Time Comparison	88
Figure (4.4): Side by side comparison between models	89
Figure (4.5): ROC curves.....	90
Figure (4.6): Model Performance Comparison	91
Figure (4.7): Pre-Language F1 Scores for Transformer Models	93
Figure (4.8). Accuracy-Speed Trade-off Across Models.....	94

List of Abbreviations

AMP	Automatic Mixed Precision
API	Application Programming Interface
AUC	Area Under the Curve
CNN	Convolutional Neural Network
CSV	Comma-Separated Values
CTkFrame	CustomTkinter Frame
DL	Deep Learning
DP	Differential Privacy
F1-score	Harmonic Mean of Precision and Recall
FPS	Frames Per Second
GPU	Graphics Processing Unit
GUI	Graphical User Interface
LSTM	Long Short-Term Memory
ML	Machine Learning
NLP	Natural Language Processing
PR-AUC	Precision-Recall Area Under Curve
ROC	Receiver Operating Characteristic
RoBERTa	Robustly Optimized BERT Pretraining Approach
RT	Retweet
SUS	System Usability Scale
TF-IDF	Term Frequency-Inverse Document Frequency
XAI	Explainable Artificial Intelligence
XLM-R	XLM-RoBERTa
ELM	Embeddings from Language Models
VADER	Valence Aware Dictionary and sEntiment Reasoner

CHAPTER ONE

INTRODUCTION TO THE STUDY

1.1 Introduction

Sentiment the subjective evaluation of ideas, events, and entities has migrated from private discourse to the global public square through Web 2.0. Social-media networks such as X (formerly Twitter), Facebook, Reddit and Instagram now broadcast billions of short messages each day, forming a rich yet volatile reservoir of public opinion that is used for strategic insight (Rodríguez-Ibáñez, Casáñez-Ventura, Castejón-Mateos, & Cuenca-Jiménez, 2023). The resulting data deluge has elevated sentiment analysis from a niche computational-linguistic task to a widely used capability in many applications, including business intelligence, public-health monitoring, and political analysis (Wankhade, Rao, & Kulkarni, 2022). Concurrently, advances in cloud computing and modern open-source machine-learning ecosystems such as HuggingFace Transformers, PyTorch, and API-based services have lowered the technical barrier to deploying sentiment-aware applications, enabling practitioners to translate social emotions into actionable dashboards in near real time (Bashiri & Naderi, 2024). Notwithstanding its appeal, extracting affect from microblog posts remains a technically challenging problem. Tweets are capped at 280 characters, rife with creative spelling, hashtags, user mentions, and rapidly evolving slang, which undermine the assumptions of traditional n-gram or lexicon methods (Wankhade et al., 2022). Moreover, users increasingly substitute words with emojis whose sentiment is whose semantic contribution varies strongly with linguistic context and multidimensional; overlooking these pictographic cues can bias polarity

estimates (Godard & Holtzman, 2022a). Sarcasm and code-switching further exacerbate ambiguity, particularly in multilingual regions such as the Gulf where Arabic and English intermingle within the same post (Abuzayed & Al-Khalifa, 2021). In response to these challenges, recent research has embraced deep-learning architectures that learn distributed representations directly from raw text. Transformer-based language models, exemplified by BERT, RoBERTa, and the domain-specific BERTweet, employ self-attention to capture long-range benefits and have established state-of-the-art performance across benchmark datasets (Bashiri & Naderi, 2024). Fine-tuning such pre-trained encoders on task-specific corpora reduces annotation costs while improving generalisation to noisy social-media vernacular (Rodríguez-Ibáñez et al., 2023). Complementary techniques data augmentation, emoji-aware tokenisation, and multilingual embedding alignment further enhance robustness to sarcasm and cross-lingual code-mixing (Abuzayed & Al-Khalifa, 2021). Despite these advances, practical deployment within the Bahraini context still faces two limitations. First, most public models are optimized for English data and underperform on the hybrid Arabic-English transcripts prevalent among local users (Wankhade et al., 2022). Second, existing academic prototypes seldom expose their capabilities through graphical interfaces, hindering adoption by non-technical decision-makers (Rodríguez-Ibáñez et al., 2023). This thesis therefore develops an end-to-end that couples a RoBERTa variant fine-tuned on a curated Twitter corpus with a custom Tkinter GUI, enabling real-time sentiment monitoring tailored to the University of Bahrain research community (Bashiri & Naderi, 2024).

The overarching aim is to design, implement, and empirically evaluate a comprehensive AI-based sentiment-analysis system capable of

reliable polarity classification across multilingual social-media streams, while also providing an integrated environment for real-time opinion polling, multilingual hashtag monitoring, model training and evaluation, transformer fine-tuning, and interactive visual analytics through an intuitive graphical interface.




Figure (1.1): Overview of AI-Based Sentiment Analysis Pipeline for Social Media.

Figure (1.1) illustrates a simplified version of the sentiment-analysis pipeline. When the user enters text, the system first performs text preprocessing and cleaning to remove unnecessary elements such as usernames, links, and other non-informative tokens. The cleaned text is then processed through the tokenization stage, which converts it into a format suitable for the sentiment classifier. The classifier generates sentiment scores based on the patterns it learned during training. Finally, these scores are passed to the application output stage, where the results are displayed in the graphical interface as charts or percentage distributions.

1.2 Research Problem

The growing reliance on social-media platforms for real-time public expression has intensified the need for systems capable of accurately interpreting sentiment from short, informal, and multilingual posts. Platforms such as X (formerly Twitter) and Reddit generate large volumes of user-generated content, yet the linguistic characteristics of this content

such as sarcasm, code-switching, emoji usage, and contextual ambiguity make automated interpretation highly challenging and frequently lead to misclassification. Traditional machine-learning classifiers, including Support Vector Machines (SVM) and Naïve Bayes (NB), also struggle with these complexities, particularly when applied to rapidly evolving, multilingual microblog environments. The methodological contribution of this research is to design and evaluate an AI-based sentiment-analysis pipeline that integrates transformer architectures, robust preprocessing methods, and multilingual tokenization strategies to overcome these limitations. The application contribution is the development of a practical, rapid, and transparent sentiment-monitoring system capable of handling real-time social-media streams within the Bahraini context. Together, these contributions address the core research problem: how to transform noisy, fast-moving social-media text into trustworthy, interpretable sentiment signals for academic, industrial, and governmental stakeholders.

1.3 Research Questions

The research questions that are investigated are:

- How accurately can the pre-trained twitter-roberta-base-sentiment model classify polarity and assign confidence scores to live Twitter streams relative to traditional sentiment-analysis approaches?
- How effectively does the CustomTkinter graphical interface enable rapid, and clear visualisation of real-time sentiment data under practical latency constraints?
- To what extent do the system's continuous opinion-polling dashboards deliver actionable insights into public sentiment

during unfolding events when compared with existing social-media analytics tools?




Figure (1.2): Research Questions on Sentiment Analysis System

1.4 Research Objectives

The objectives of this research are as follows:

- Develop an end-to-end, real-time sentiment-analysis system that streams Twitter data, preprocesses it, and classifies polarity and confidence scores using the pre-trained transformer model `twitter-roberta-base-sentiment`.
- Benchmark the model's classification accuracy, latency, and resource efficiency against rule-based and classical machine-learning baselines to quantify performance improvements.
- Design and implement a CustomTkinter graphical interface that presents live polarity distributions, confidence metrics, and temporal sentiment trends through intuitive, interactive visualisations.
- Implement and validate an opinion-polling module that aggregates classified tweets into continuous public-sentiment indicators, assessing its utility for timely, actionable insight generation.

1.5 Research Significance

The importance and relevance of this research study can be attributed to the following factors:

- Advances the academic discourse on social-media analytics by demonstrating the efficacy of domain-specific transformers for real-time sentiment classification, thus providing a replicable benchmark for future AI research.
- Supplies practitioners, policymakers, and marketers with a practical mechanism to monitor and interpret public opinion dynamics on Twitter, enhancing situational awareness during rapidly evolving events.
- Facilitates data-driven decision-making by delivering immediate polarity and confidence metrics, enabling timely adjustments to communication strategies, crisis responses, and campaign messaging.
- Showcases an integrated, user-centric platform that unifies streaming data ingestion, transformer-based analysis, graphical visualization, and opinion polling, thereby illustrating a scalable template for interactive sentiment-analysis applications.

1.6 Study Setting

The study was conducted through University of Bahrain using a Windows 11 workstation equipped with an Intel Core i7 CPU, 32 GB RAM, and an NVIDIA RTX 3060 Ti GPU. Python 3.10 served as the primary language, with the Hugging Face Transformers library employed to load the “twitter-roberta-base-sentiment” model. Tweet streams were simulated through Tweepy’s API sandbox to ensure continuous data flow

during testing. A CustomTkinter-based GUI was developed and executed within a Jupyter Notebook and later packaged as a standalone desktop application.

1.7 Research Terminologies

Sentiment polarity refers to the categorical orientation positive, negative, or neutral assigned to each tweet according to its expressed attitude toward a target entity. Confidence score denotes the probabilistic certainty, scaled between 0 and 1, that the underlying classifier attaches to the predicted polarity, thereby guiding thresholding decisions. A pre-trained model is a neural network whose parameters have been optimised on a large external corpus and subsequently reused or fine-tuned for the specific task, drastically reducing data and computation requirements. Transformer-based model designates any architecture utilising multi-head self-attention to capture contextual dependencies within sequences; `twitter-roberta-base-sentiment` is the specialised transformer employed in this study. Tweet preprocessing encompasses the automated cleansing and normalisation steps such as lower-casing, tokenisation, emoji translation, stop-word removal, and URL masking executed before model inference to mitigate noise. Opinion polling describes the temporal aggregation of individual tweet classifications into rolling statistics that approximate collective public sentiment on a defined topic or hashtag. GUI output represents the real-time visual elements rendered by the CustomTkinter interface, including colour-coded polarity bars, time-series plots of aggregated sentiment, and textual confidence indicators delivered to end users. Streaming pipeline denotes the asynchronous collection, buffering, and dispatch of tweets from the Twitter API to the classification engine, ensuring minimal latency between data

acquisition and visual presentation under typical broadband network conditions. Together, these terminologies form the conceptual framework guiding system design, experimental evaluation, and result interpretation throughout the thesis.

CHAPTER TWO

LITERATURE REVIEW

2.8 Overview of Sentiment Analysis

2.8.1 Historical Development of Sentiment Analysis

Sentiment analysis has evolved through several methodological phases, beginning with rule-based systems, progressing to machine-learning classifiers, and eventually advancing to deep learning and transformer-based models. Following Figure 2.1. Early rule-based approaches relied on hand-crafted lexicons and grammatical rules to assign polarity, but they often failed to capture contextual nuances such as negation, sarcasm, and informal language variations common in social media. The introduction of machine-learning approaches marked a significant shift. Algorithms such as Naïve Bayes, Support Vector Machines, and Logistic Regression leveraged labeled datasets to learn sentiment patterns statistically. These models outperformed lexicon methods but still depended heavily on manual feature engineering. The emergence of deep learning brought substantial improvements. CNNs captured local n-gram features, whereas LSTMs and other recurrent networks modelled long-range dependencies, enabling more robust understanding of informal and noisy text.




Figure (2.1): Historical Development of Sentiment Analysis

The most recent phase is dominated by transformer-based architectures, particularly BERT and RoBERTa, which use self-attention mechanisms to capture global context and deliver state-of-the-art performance across multiple sentiment benchmarks, including multilingual and social-media datasets. Despite remarkable progress from naïve lexicon lookups to sophisticated AI models, challenges remain. Current state-of-the-art models struggle with sarcasm detection, idioms, emerging slang, and domain shifts. They often function as "black boxes" with limited interpretability. Future research explores hybrid approaches combining language models with external knowledge, multimodal sentiment analysis integrating text with audio-visual cues, and developing more explainable and fair models addressing training biases.

2.8.2 Granularity Levels and Applications of Sentiment Analysis

Sentiment analysis can be performed at different levels of granularity, each serving distinct analytical goals. At the document level, entire texts such as product or movie reviews receive a single sentiment label, providing a coarse but efficient overview of overall opinion (Su et al., 2023). Sentence-level analysis evaluates individual sentences, making it especially suitable for short-form content like tweets, where a single sentence often conveys a complete sentiment. Aspect-based sentiment analysis (ABSA) offers the most fine-grained approach by identifying specific aspects within a text and assigning sentiment individually. For instance, in the review “*The camera is great but the battery is awful*,” the same sentence expresses contrasting sentiments toward different features. SemEval datasets such as SemEval-2014 Task 4 on restaurant and laptop reviews have established standardized corpora and evaluation protocols for ABSA (Pontiki et al., 2016). These relationships between granularity levels and their

applications are illustrated in Figure 2.2. Across these levels, social media introduces additional challenges. Tweets are short, informal, and often contain emojis, abbreviations, sarcasm, and mixed-language expressions. Sarcasm can invert the intended sentiment of positive words, negation can subtly shift polarity (e.g., “not great”), and code-mixed or multilingual posts complicate analysis when models are trained primarily on English (Helal et al., 2024; Yacoub et al., 2024). These factors make Twitter an especially demanding testbed for sentiment analysis. Despite these challenges, sentiment analysis has been widely applied in finance (market reaction monitoring), healthcare (tracking mental health indicators), e-commerce and marketing (brand monitoring and product review mining), and political analysis (public opinion tracking during elections) (Greyling & Rossouw, 2023; Helal et al., 2024; Yacoub et al., 2024; Sudirjo et al., 2023). In all these domains, short real-time messages from platforms such as Twitter and Reddit play a central role.




Figure (2.2):Granularity levels of sentiment analysis and their associated real-world applications.

This thesis focuses specifically on sentence-level and tweet-level social media sentiment analysis, within an environment characterized by short, noisy, and mixed Arabic–English posts. The aim is to leverage domain-adapted RoBERTa/XLM-R models to provide accurate and practical sentiment analysis through an accessible GUI.

2.8.3 Challenges in Social Media Sentiment Analysis

Social media sentiment analysis faces several unique challenges that stem from the informal linguistic style of user-generated content and the technical constraints of online platforms. Linguistically, social media language is highly informal and creative. Posts often include slang, abbreviations, and code switching, which complicate text processing (Indumathi, 2020). Mining sentiments from social data introduces concerns about privacy, misinformation, and authenticity. User privacy is a paramount issue even if posts are public, they may contain sensitive personal data, and users often do not expect their content to be mined a masse. Analysing social posts can lead to privacy violations if proper care isn't taken (Vogel & Lange, 2023). Researchers must adhere to data protection regulations and ethical guidelines, e.g. anonymizing user identifiers. In summary, social media sentiment analysis must overcome noisy, short, and idiosyncratic text; navigate platform-imposed data access limits and real-time processing demands; and ensure ethical use of data amidst privacy concerns, misinformation, and inauthentic actors. Ongoing research is addressing these challenges with more sophisticated language models (e.g., sarcasm- and code-switching-aware algorithms), robust big-data infrastructures, and privacy-conscious data handling methods.

2.9 Traditional Approaches to Sentiment Analysis

2.9.1 Lexicon-Based Methods

Lexicon-based sentiment analysis relies on predefined dictionaries that assign polarity scores to words or phrases. These methods compute overall sentiment by aggregating the scores of words appearing in a text. Their main strengths are simplicity, transparency, and the absence of training requirements, making them suitable for quick baseline evaluations. Several well-known English lexicons have been widely used in social media research. VADER (Valence Aware Dictionary and sEntiment Reasoner) is specifically designed for online text and incorporates heuristics for capitalization, punctuation, negation, and emojis, making it more effective on tweets than generic lexicons (Hutto & Gilbert, 2014). AFINN (Finn Årup Nielsen) provides integer-based polarity scores and covers informal language, while SentiWordNet extends sentiment scoring to WordNet, although it lacks coverage for slang and non-standard expressions common on Twitter. Multilingual sentiment lexicons have been created to extend analysis beyond English. Resources such as ML-SentiCon and the NRC Emotion Lexicon provide polarity or emotion categories across multiple languages, while SenticNet maps sentiment at a concept level to support broader semantic interpretation. However, these lexicons still struggle with dialectal variations, transliteration (e.g., Romanized Arabic), and rapidly evolving slang that characterize social media platforms. Despite their advantages, lexicon-based methods face notable limitations. They provide no contextual understanding, cannot reliably interpret sarcasm, idioms, or complex negation, and offer limited support for mixed-language content. These limitations reduce their effectiveness on noisy, code-mixed Arabic—

English tweets. As a result, lexicon approaches are often used only as baselines, with contemporary systems favoring machine-learning and transformer-based models that can better handle informal and multilingual social media text.

2.9.2 Machine Learning Approaches

Classical machine learning methods played a central role in early sentiment analysis and continue to serve as strong baselines for short-text classification tasks, including social media analysis. These models rely on numerical feature representations such as Bag-of-Words (BoW), TF-IDF, and n-grams, transforming text into sparse vectors before applying statistical classifiers. Naïve Bayes (NB) models estimate the probability of a sentiment class based on word occurrences under an independence assumption. NB is computationally efficient and performs well on high-dimensional sparse data but struggles with contextual dependencies.

Support Vector Machines (SVM) became widely used for sentiment classification due to their effectiveness with high-dimensional text features. SVMs learn a decision boundary that maximizes class separation and often outperform simpler classifiers on short, informal texts such as tweets. Also, Logistic Regression (LR) provides a probabilistic framework for binary or multiclass sentiment prediction. It is known for its interpretability and competitive performance, especially with regularized feature weights. Decision Tree (DT) models offer intuitive, rule-based representations but tend to overfit due to the complexity and sparsity of textual data. Random Forests (RF) alleviate this through ensemble learning, improving generalization by aggregating predictions from multiple trees. However, both DT and RF require more computation than

linear models and provide less interpretability than LR or NB. Beyond supervised learning, semi-supervised methods such as self-training and active learning help address the scarcity of labeled data in social media contexts. These techniques iteratively incorporate model-predicted labels or selectively query human annotators to expand training sets efficiently. Unsupervised clustering methods like K-means can reveal latent sentiment groupings within unlabeled corpora, supporting exploratory analysis or pre-labeling.

Overall, classical machine learning methods are fast, interpretable, and effective on structured or moderately noisy text. However, they lack contextual understanding and struggle with informal, multilingual, or code-mixed content. These limitations make them less suitable for Twitter sentiment analysis—particularly for mixed Arabic–English data—where transformer-based models now provide substantially stronger performance.

2.9.3 Feature Engineering for Sentiment Analysis

Feature engineering transforms raw social media text into numeric representations, critical for traditional machine learning (ML)-based sentiment analysis. These engineered features broadly fall into traditional text, linguistic, and statistical/semantic categories. Traditional text features are foundational and include Bag-of-Words (BoW), Term Frequency-Inverse Document Frequency (TF-IDF), and n-grams (word and character). BoW features count word occurrences without considering order, making them efficient yet simplistic. Naïve Bayes classifiers notably perform effectively with binary BoW vectors due to independence assumptions (Kumar et al., 2025).

Statistical and semantic features embed deeper linguistic and domain-specific knowledge into sentiment models. Pre-trained word embeddings (Word2Vec, GloVe) represent semantic similarity as dense vectors, enhancing the generalization capabilities of classifiers by capturing relationships between sentiment words (Kumar et al., 2025). Lexicon-based sentiment scores from resources like VADER or AFINN directly quantify word sentiment, enriching ML models with external knowledge. Comparatively, traditional features like TF-IDF are advantageous for SVMs due to their sparsity, while Naïve Bayes excels with simple BoW. Logistic Regression leverages lexicon-weighted features effectively, combining interpretability with improved accuracy. A typical sentiment analysis pipeline incorporates these varied features, as

Deep Learning in Sentiment Analysis

2.9.4 Recurrent Neural Networks (LSTM/GRU)

Recurrent Neural Networks (RNNs) are specialized for processing sequential data, where each step depends on previous outputs. In sentiment analysis, they help capture contextual nuances across time. However, vanilla RNNs face challenges due to vanishing and exploding gradients. These issues limit their ability to model long-term dependencies effectively. As the sequence length increases, the learning signal deteriorates, leading to poor performance on tasks requiring memory over many timesteps.

To address these limitations, Long Short-Term Memory (LSTM) networks were introduced. LSTMs enhance RNNs by incorporating a memory cell and three gating mechanisms: the forget gate, input gate, and output gate. These gates manage the flow of information, allowing the model to retain relevant information and discard irrelevant content across

sequences. This architecture enables the LSTM to retain or discard past information selectively, thus addressing the vanishing gradient issue. LSTMs are well-suited for sentiment analysis, especially when dealing with long texts or complex sentence structures, as they maintain context across multiple time steps (Waqas & Humphries, 2024). Gated Recurrent Units (GRUs) offer a streamlined alternative to LSTMs. They simplify the architecture by merging the forget and input gates into a single update gate and removing the cell state. The GRU uses two main gates update and reset to regulate memory. GRUs are computationally efficient and perform comparably to LSTMs, particularly in environments with limited data or where faster training is necessary. Their reduced parameter count helps in avoiding overfitting while enabling robust performance (Rivas et al., 2025). In comparing the two, LSTMs are more expressive and suitable for modeling complex dependencies in long sequences. Their distinct memory cell helps in tasks involving nuanced sentiment changes over time. However, GRUs train faster and generalize better on smaller datasets due to fewer parameters. A study by (Raza et al. 2021) revealed that GRUs achieved higher accuracy than LSTMs and RNNs on cloud service reviews, highlighting their efficiency and robustness in resource-constrained settings. Numerous applications demonstrate the efficacy of LSTM and GRU architectures in sentiment analysis. (Nistor et al. 2021) applied LSTMs to Twitter sentiment classification and achieved over 82% accuracy, outperforming traditional ML models like SVM. (Alosaimi et al. 2024) combined AraBERT with LSTM to create a robust Arabic sentiment classifier that surpassed 97% accuracy. These examples illustrate LSTM's strength in handling morphologically rich and long textual data.

GRUs have also been widely adopted in multilingual sentiment tasks, especially where rapid training and data efficiency are needed. For

instance, studies leveraging bilingual word embeddings and GRUs have successfully classified sentiments in code-mixed tweets, showing the model's adaptability across languages and informal language styles. Hybrid models combining LSTM or GRU layers with CNN or transformer-based encoders have shown superior performance. (Bouassida and Mezali 2025) integrated RoBERTa with a BiLSTM and reported accuracy above 91% on IMDb and Twitter datasets. These architectures capture both global context from transformer encoders and sequential dynamics from LSTM/GRU layers, enhancing sentiment prediction across domains. In conclusion, both LSTM and GRU networks offer solutions to the limitations of traditional RNNs. LSTMs excel in modeling long-term dependencies with rich internal state control, while GRUs provide speed and simplicity with comparable accuracy. Their integration into modern sentiment analysis systems either as standalone classifiers or in hybrid forms continues to advance the field, powering real-world applications from social media monitoring to multilingual opinion mining.

2.9.5 Convolutional Neural Networks for Text

Convolutional Neural Networks (CNNs) have been successfully adapted from computer vision to text classification tasks, offering efficient local pattern recognition capabilities. A typical CNN for sentiment analysis begins by mapping each word to dense embedding vectors, creating a 2D matrix of shape $n \times d$. One-dimensional convolutional filters slide over these embeddings to extract local n -gram features, followed by max-pooling and a softmax classification layer. Multi-channel CNNs enhance feature diversity by processing multiple input representations in parallel. Examples include architectures using static vs. trainable embeddings (Kim,

2014), multi-scale filters (Rajesh & Hiwarkar, 2023), and channels combining embeddings with syntactic or metadata features.

Table (2.1): CNN comparisons

Model Type	Key Idea	Example Studies
Single-Channel CNN	Uses one embedding representation with convolution + pooling.	(Kim 2014)
Multi-Channel CNN (Static +Trainable)	Combines fixed pre-trained embeddings with trainable embeddings.	(Kim 2014)
Multi-Scale CNN	Parallel convolutions with filter sizes 3, 4, 5, 6, 8.	(Rajesh & Hiwarkar 2023)
Feature-Fusion CNN	Combine embeddings with syntax/positional features.	(Feng & Cheng 2021)
Metadata-Augmented CNN	Parallel channels for text + Twitter metadata.	(Sitaula & Shahi 2022)

2.9.6 Word Embeddings and Contextual Representations

Word embeddings are a core component of modern sentiment analysis, enabling models to represent words as continuous vectors rather than isolated tokens. Early approaches used static embeddings, where each word type received one fixed vector regardless of its context. Models such as Word2Vec (Mikolov et al., 2013), GloVe (Pennington et al., 2014), and FastText (Bojanowski et al., 2017) fall into this category. Word2Vec learns representations based on local word co-occurrences, GloVe incorporates global statistical information, while FastText improves robustness by representing words through character-level n-grams. FastText is particularly effective for noisy social media text because it can generate vectors for misspellings, slang, and unseen words. However, static embeddings cannot adapt to different meanings of the same word. A word like “*apple*” receives the same vector whether it refers to fruit or the technology company. This limitation creates challenges for sentiment interpretation, especially in short, ambiguous social media posts where

context is essential. The shift toward contextual embeddings transformed sentiment analysis. However, transformer-based architecture like BERT, RoBERTa, and XLM-R, generate word representations that depend on the surrounding sentence. These contextual models capture polysemy, idiomatic expressions, and long-range dependencies far more effectively than static embeddings. Their ability to embed meaning dynamically makes them particularly suitable for noisy platforms like Twitter. For multilingual applications, XLM-RoBERTa provides strong cross-lingual semantic alignment across 100 languages, enabling models trained on English sentiment data to generalize to other languages, including dialectal Arabic and mixed Arabic–English tweets. This capability is crucial for platforms where code-switching and transliterated forms (e.g., Romanized Arabic) are widespread. Empirical studies consistently show that contextual embeddings outperform static vectors on sentiment classification benchmarks. While Word2Vec or FastText combined with CNN/LSTM architecture provides solid baselines, transformer-based models achieve substantially higher accuracy, especially when fine-tuned on domain-specific datasets such as TweetEval. Given these advantages, this thesis adopts RoBERTa and XLM-R as the primary embedding and encoding models. Their contextual, multilingual, and domain-adaptive properties make them far more suitable for analyzing noisy, mixed Arabic–English Twitter content than static embedding methods.

2.10 Transformer Architectures

2.10.1 Attention Mechanism and Self-Attention

Attention mechanisms revolutionized neural network architectures by enabling models to selectively focus on relevant input components.

Traditional encoder-decoder models employ attention to align and weight encoder hidden states during output generation. Bahdanau's additive attention computes alignment scores between decoder state s_{t-1} and encoder output h_i through feed-forward networks producing attention weights summing to unity. Context vectors $c_t = \sum_i \alpha_{t,i} h_i$ form weighted sums of encoder states, enabling decoders to attend to relevant input portions at each step through global soft-attention mechanisms (Kardakis et al., 2021). The attention process follows systematic steps: for each decoder timestep, scores are computed for all encoder states using either additive or multiplicative formulations, converted to probabilities via SoftMax, and combined into context vectors through weighted summation. This dynamic alignment between target and source sequences enables flexible correspondence modeling, with Bahdanau's additive attention capturing more complex alignments while Luong's dot-product approach operates faster. Self-attention represents a paradigmatic shift from cross-attention by relating tokens within identical sequences. Each token generates query, key, and value vectors (Q, K, V) through linear embedding projections, enabling every token to attend to all others simultaneously. The scaling factor $\sqrt{d_k}$ prevents extreme softmax weights, while inherent parallelizability across positions eliminates sequential processing requirements. This mechanism replaces RNN functionality by allowing each position to incorporate information from all other positions within single layers, facilitating faster training and more comprehensive context modeling (Conneau et al., 2020). Modern implementations utilize pre-trained transformer models like RoBERTa and XLM-RoBERTa, comprising stacked transformer layers with multi-head self-attention. XLM-RoBERTa exemplifies multilingual capabilities through pre-training on 2.5TB CommonCrawl text spanning 100 languages. The base

architecture features 12 layers with 12 attention heads per layer (768-dimensional model). During processing, special [CLS] tokens are prepended to input sequences, with final [CLS] hidden states serving as pooled representations for classification. The `xlm_roberta_analyzer` module leverages these [CLS] vectors through feed-forward classifier heads for sentiment prediction (Prytula, 2024).

Multi-head self-attention extends basic mechanisms by running parallel attention functions on different learned linear projections. Each head projects Q, K, V into distinct subspaces, computing individual attention maps before concatenation and linear back-projection. This design enables simultaneous capture of multiple relationship types or features. As Ni et al. (2023) explain, multi-head attention allows models to attend to "information from different representation subspaces" concurrently. Key distinctions emerge between attention types: classical Bahdanau/Luong attention operates in encoder-decoder architectures through sequential alignment weight computation, conditioning each output token on source-dependent contexts. Self-attention functions within single sequences, enabling simultaneous attention across all positions through parallel computation. While RNN-based attention provides sequential source-to-target alignment, transformer self-attention offers flexible, parallelizable contextual information integration across sequences through multiple representation subspaces. These mechanisms fundamentally transformed sentiment analysis by enabling more sophisticated context modeling, long-range dependency capture, and efficient parallel processing. The evolution from sequential cross-attention to parallel self-attention represents a critical advancement in neural architecture design, particularly beneficial for social media text analysis

where context and nuanced relationships significantly impact sentiment interpretation.

2.10.2 BERT and Its Architecture

Bidirectional Encoder Representations from Transformers (BERT) represents a revolutionary deep, encoder-only Transformer model that fundamentally changed natural language processing. BERT's input representation combines three embedding types for each token: WordPiece token embeddings, segment embeddings distinguishing sentence boundaries, and positional embeddings encoding token indices (Gardazi et al., 2025). BERT-base comprises 12 identical Transformer encoder layers with 768 hidden dimensions, 12 attention heads, and approximately 110 million parameters. Each BERT layer contains two sublayers with residual connections and layer normalization. The multi-head self-attention mechanism enables each token to attend simultaneously to all sequence tokens, capturing both left and right contexts bidirectionally. The position-wise feed-forward network applies two-layer perceptions to each position with intermediate non-linearity. Residual connections and layer normalization stabilize training and facilitate gradient flow through the 12-layer architecture.

Implementation involves 80% token replacement with [MASK], 10% random token substitution, and 10% unchanged tokens to mitigate train-test discrepancies (Wettig et al., 2023). Next Sentence Prediction (NSP) constitutes the second pretraining objective, requiring binary classification of sentence pairs to determine sequential coherence. The loss trains BERT to capture inter-sentence relationships, though subsequent research questioned its effectiveness, leading models like RoBERTa to abandon this objective. BERT's architectural innovations distinguish it

from related models. Compared to ELMo's (Embeddings from Language Models) shallow bidirectional LSTM producing surface-level bidirectional context at final layers, BERT implements deep bidirectional attention at every layer through self-attention mechanisms. This deeper contextual integration yields substantially richer semantic representations. Several BERT variants address specific limitations while preserving core architectural principles. DistilBERT achieves 40% parameter reduction (66M versus 110M) through knowledge distillation, maintaining 96% of BERT's performance while operating 45-60% faster during inference (Barbon & Akabane, 2022). ALBERT introduces factorized embeddings and cross-layer parameter sharing, reducing parameters from $O(V \times H)$ to $O(V \times E + E \times H)$ where embedding dimension E is smaller than hidden dimension H , achieving only 12M parameters while maintaining comparable performance (Zeng et al., 2022).

RoBERTa preserves BERT's encoder architecture while eliminating NSP and utilizing enhanced training procedures with larger datasets and improved masking strategies. XLM-RoBERTa extends this approach to 100 languages through multilingual masked language pretraining, achieving state-of-the-art cross-lingual performance by scaling BERT's concepts globally (Conneau et al., 2020). Modern sentiment analysis systems leverage RoBERTa and XLM-R as underlying language models, benefiting from BERT's architectural innovations while incorporating training improvements. These encoder-only models provide robust contextual representations that significantly enhance sentiment classification accuracy across diverse languages and domains, establishing BERT's architecture as the foundation for contemporary transformer-based sentiment analysis applications.

2.10.3 RoBERTa and Optimization Techniques

RoBERTa (Robustly Optimized BERT Approach) retains BERT's encoder-only Transformer architecture and bidirectional attention mechanisms while implementing substantial pretraining regime modifications for enhanced performance (Wang et al., 2025). Rather than architectural changes, RoBERTa achieves superior results through systematic training optimizations, demonstrating that careful pretraining procedures can significantly impact model effectiveness without structural modifications.




Figure (2.3): BERT-large Architecture (Dikmener, 2022)

The most significant modification involves eliminating BERT's Next Sentence Prediction (NSP) objective. Original NSP aimed to help models learn sentence-pair relationships, but empirical evidence revealed minimal contribution to most downstream tasks. By removing NSP, RoBERTa concentrates all pretraining effort on masked token prediction, improving token-level contextual understanding while avoiding capacity

waste on irrelevant sentence-pair tasks for many applications. Dynamic masking represents another crucial optimization. Unlike BERT's fixed masking patterns where identical tokens are masked repeatedly, RoBERTa implements dynamic masking by re-sampling masked tokens randomly for each training epoch and sequence. This approach exposes models to diverse masking configurations continuously, preventing overfitting to specific mask patterns while enabling learning from richer token-mask contexts, ultimately improving generalization capabilities.

RoBERTa's training corpus expansion dramatically exceeds BERT's scale, utilizing over 160GB of uncompressed text from multiple sources including Wikipedia, BookCorpus, CommonCrawl News, OpenWebText, and Stories datasets, compared to BERT's approximately 16GB. This order-of-magnitude increase provides exposure to substantially wider language patterns and vocabulary coverage, with experiments consistently demonstrating superior performance from larger training corpora. Optimization parameter scaling represents another key advancement. RoBERTa employs batch sizes reaching thousands of sequences (up to 8,000 sentences per batch) with correspondingly adjusted learning rates. Larger batches enable processing more training examples per weight update, stabilizing gradient estimates when learning on enormous datasets. By aggregating over numerous examples, models can train for extended iterations without overfitting, with appropriately tuned learning rate schedules ensuring steady convergence at scale. RoBERTa utilizes maximum-length inputs (512 tokens) throughout pretraining, contrasting with BERT's approach of initially using shorter sequences before gradually increasing length. This full-length training from the start enables better modeling of longer-range language dependencies across extended passages. Additionally, RoBERTa implements large byte-pair encoding

(BPE) vocabulary (~50,000 tokens) similar to GPT-2, significantly exceeding BERT's ~30,000 WordPieces. This expanded vocabulary handles out-of-vocabulary words more gracefully through subword representation. These modifications yield measurable performance improvements. Original RoBERTa experiments achieved state-of-the-art GLUE benchmark results with scores reaching 88.5, surpassing corresponding BERT checkpoints through pretraining optimizations alone. On social media benchmarks, RoBERTa consistently outperforms vanilla BERT. Barbieri et al. (2020) demonstrated through TweetEval evaluations that pretrained RoBERTa-base models frequently exceeded Twitter-specific trained models, indicating robust pretraining's effectiveness for social media text processing. Contemporary implementations leverage specialized variants like CardiffNLP's "twitter-roberta-base-sentiment" model, combining RoBERTa's optimization advantages with domain-specific training. This model, pretrained on tens of millions of English tweets from January 2018 to December 2021 and fine-tuned using TweetEval benchmarks, exemplifies how RoBERTa's robust pretraining foundation enhances domain-specific applications (Semary et al., 2024). Such specialized models demonstrate RoBERTa's versatility for handling short, informal social media text while maintaining the architectural simplicity and training efficiency improvements that distinguish it from BERT's original implementation.

2.10.4 XLM-RoBERTa Overview

XLM-RoBERTa (XLM-R) represents a massively multilingual extension of RoBERTa, preserving the same encoder-only Transformer architecture while dramatically expanding cross-lingual capabilities. Trained as a masked language model on the extremely large

CommonCrawl-derived CC100 corpus covering approximately 100 languages, XLM-R utilizes over 2.5 terabytes of filtered text to learn robust cross-lingual representations from unlabeled data (Conneau et al., 2020). This massive scale enables XLM-R to acquire strong zero-shot transfer abilities, with Conneau et al. reporting that XLM-R "significantly outperforms multilingual BERT (mBERT) on a variety of cross-lingual benchmarks, including +14.6% average accuracy on XNLI."

XLM-R's design incorporates several key innovations that distinguish it from previous multilingual models. The CC100 dataset construction from year-2018 web crawls across 100+ languages provides unprecedented multilingual coverage, exceeding earlier multilingual corpora by orders of magnitude. This massive data scale and broad language representation prove critical to XLM-R's superior cross-lingual encoding capabilities. The training methodology extends standard masked language modeling objectives across all 100 languages in CC100, forcing the model to align representations across diverse linguistic families. This unified multilingual training creates a common semantic space where English sentiment data can effectively predict sentiment in other languages, enabling robust zero-shot and few-shot transfer capabilities. Empirical evidence demonstrates substantial transfer gains through multilingual MLM training, as evidenced by significant XNLI performance improvements.

XLM-R implements SentencePiece unigram language model tokenization with approximately 250,000 subword tokens trained on CC100, contrasting with mBERT's WordPiece vocabulary of ~119,000 tokens. This larger SentencePiece vocabulary enables more comprehensive coverage of language-specific subwords and scripts, though potentially resulting in longer sequences. This design choice facilitates more even

modeling across typologically diverse languages by ensuring each language contributes substantially to joint vocabulary. Comparative evaluations consistently demonstrate XLM-R's superiority over other multilingual encoders. Against mBERT baselines, XLM-R's larger training set and capacity translate into substantially stronger transfer performance, with 14.6% higher average accuracy on cross-lingual NLI tasks. In multilingual TweetEval benchmarks, XLM-R variants significantly outperform mBERT on tweet classification tasks, with double-digit margin improvements on tasks like MLQA and NER (Barbieri et al., 2022).

Recent multilingual alternatives include Google's RemBERT and mT5, though XLM-R maintains practical advantages for sentiment analysis applications. RemBERT achieves comparable performance through decoupled input/output embeddings and parameter efficiency improvements, while mT5's sequence-to-sequence architecture provides state-of-the-art multilingual performance at higher computational costs unsuitable for encoder-focused pipelines. Empirical studies consistently validate XLM-R's effectiveness for multilingual sentiment analysis. Xu et al. (2022) note that "XLM-Roberta was pre-trained in 100 languages and obtained state-of-the-art performance on cross-lingual classification, sequence labeling and question answering, especially in low-resource languages." Twitter-specific evaluations demonstrate that XLM-R variants fine-tuned on social media data typically yield strongest performance for multilingual sentiment tasks. Contemporary implementations utilized specialized models like CardiffNLP's "twitter-xlm-roberta-base-sentiment-multilingual," which combines XLM-R's foundational cross-lingual capabilities with domain-specific Twitter training across thirty languages. This model exemplifies how XLM-R's robust multilingual foundations enable effective real-time sentiment classification across diverse languages

while maintaining the architectural efficiency and transfer learning advantages that distinguish it from earlier multilingual approaches. Such implementations demonstrate XLM-R's practical value for multilingual social media sentiment analysis applications requiring both cross-lingual transfer capabilities and domain-specific performance optimization.

2.11 Fine-Tuning Strategies for Sentiment Analysis

2.11.1 Evolution of Transfer Learning Approaches

The emergence of pre-trained language models has reshaped sentiment analysis, with fine-tuning becoming essential for achieving state-of-the-art results. Early transfer-learning methods treated models such as ELMo and BERT as frozen feature extractors, training only a task-specific classifier on top (Peters et al., 2018; Devlin et al., 2019). This frozen-model strategy remains highly efficient, as no transformer parameters are updated, which significantly reduces memory usage and training cost. Kamruzzaman and Kim (2023), for example, demonstrate that frozen RoBERTa embeddings combined with FastText or SVM classifiers can produce competitive results with minimal resources. However, frozen representations offer limited adaptability to the nuances of specific sentimental tasks. Full fine-tuning represents the opposite approach, updating all parameters of the pre-trained model. This yields strong gains especially with RoBERTa and XLM-R by allowing the model to specialize in sentiment patterns found in social-media text. The drawback, however, is the substantial computational demand, motivating the development of more efficient intermediate techniques. Partial fine-tuning strategies aim to balance accuracy and efficiency. Adapter-based methods (Houlsby et al., 2019) freeze the full model and insert lightweight trainable modules within each transformer

block, achieving near-full fine-tuning performance with far fewer parameters. Selective-layer training provides another compromise: Lee et al. (2019) show that fine-tuning only the top quarter of the transformer layers can recover about 90% of full-model accuracy, reflecting the division between lower-layer general linguistic features and upper-layer task-specific features.

These tuning strategies have become especially effective in social-media sentiment analysis. The TweetEval benchmark (Barbieri et al., 2020) illustrates that domain-specific adaptation consistently outperforms generic fine-tuning. CardiffNLP’s twitter-roberta-base-sentiment model exemplifies this pattern by undergoing extra pre-training on 124 million tweets before fine-tuning, achieving superior performance over standard RoBERTa. Similarly, XLM-T (Barbieri et al., 2022) extends XLM-R through additional multilingual Twitter pre-training, outperforming general-domain multilingual models across 30+ languages. Collectively, findings across these studies confirm that modern systems typically fine-tune RoBERTa for English and XLM-R for multilingual contexts using either full fine-tuning or parameter-efficient methods establishing them as core components of high-accuracy sentiment analysis on social-media platforms.

2.11.2 Domain Adaptation Techniques

The informal and highly variable nature of social media text creates substantial domain shift challenges that significantly hinder cross-domain sentiment analysis. Platforms such as Twitter and Reddit contain linguistic patterns—slang, abbreviations, emojis, misspellings, hashtags, and code-switching—that differ sharply from formal domains like news or reviews.

Barbieri et al. (2022) describe Twitter as an “uncurated” environment rich in misspellings, slang, vulgar expressions, and multimodality, which disrupts traditional linguistic signals and causes models trained on formal datasets to generalize poorly without domain adaptation. Adversarial domain adaptation has become a leading solution to this problem. Methods based on Domain-Adversarial Neural Networks (Ganin et al., 2016) train models to produce domain-invariant representations by jointly optimizing sentiment accuracy and resistance to domain discrimination. Recent work by Lu et al. (2023) shows that adversarial modules can effectively align feature distributions between source text (e.g., reviews) and target social-media text, improving cross-domain performance. In practice, these systems often use labeled data from formal domains combined with unlabeled Twitter data, enabling the encoder to learn sentiment features that remain stable across domain boundaries. Self-training and pseudo-labeling offer another effective strategy for leveraging unlabeled social-media data. A model trained on source-domain labels generates predictions on unlabeled tweets; high-confidence predictions are then used as pseudo-labels in iterative cycles. Lu et al. (2023) further enhance this approach with a meta-learning-based framework (DaMSTF) that jointly optimizes model parameters and pseudo-labels, enabling efficient transfer of sentiment knowledge to informal domains. This method is especially effective when large volumes of unlabeled tweet data are available.

Intermediate-task tuning, or continued pre-training, represents a third major adaptation technique. This involves further masked language model training on target-domain text before sentiment fine-tuning. Idrissi-Yaghir et al. (2023) demonstrate this by continuing pre-training on 1.2 million German transit-related tweets, significantly improving downstream

accuracy. Large-scale implementations such as XLM-T follow this approach by continuing XLM-R’s pre-training on millions of tweets across multiple languages before task-specific fine-tuning. Related strategies also include pre-training on emoji prediction or platform-specific tasks, which help the model acquire Twitter-style linguistic cues. Together, these methods adversarial feature alignment, pseudo-labeling, and continued in-domain pre-training enable transformer models like RoBERTa and XLM-R to adapt effectively to Twitter’s informal, noisy language. Empirical results consistently show that Twitter-specific models such as XLM-T outperform general-domain versions, underscoring the importance of robust domain adaptation in social-media sentiment analysis.

2.11.3 Few-Shot and Zero-Shot Learning

In resource-constrained scenarios involving new languages or domains without annotated sentiment data, few-shot and zero-shot learning methods provide attractive alternatives to traditional supervised approaches. These techniques become particularly valuable when developing sentiment analysis systems for low-resource languages or specialized domains where labeled data acquisition is challenging or expensive. Prompt-based techniques leverage large pre-trained language models such as GPT variants and T5 models by framing sentiment analysis as text completion or question-answering tasks. Zero-shot approaches provide no labeled examples in the prompt, while few-shot methods include several labeled tweet examples as demonstrations to guide the model's predictions. These approaches exploit the broad world knowledge encoded in large language models, enabling sentiment inference even without task-specific fine-tuning. Prior research demonstrates that GPT-3 can perform impressively across many tasks in few-shot settings, though

sentiment analysis results remain mixed, with large models sometimes underperforming compared to fine-tuned smaller models, particularly in narrow domains requiring precise calibration.

Despite performance variability, prompt-based methods have proven valuable in practical applications. Few-shot GPT prompting has been successfully employed to generate synthetic code-mixed tweets for data augmentation, indirectly improving sentiment classifiers trained on these synthetic examples (Zeng, 2024). Moreover, (Kumar and Albuquerque 2021) demonstrated this approach by fine-tuning XLM-R on English data and achieving 60.9% accuracy on Hindi tweets through zero-shot transfer. Similar work shows XLM-R achieving 50-75% accuracy on binary sentiment classification for languages like Swahili and Yoruba without target language training data. The comparative advantages of these approaches reflect different use case priorities. Prompt-based few-shot methods with GPT or Flan-T5 offer maximum flexibility and often capture nuanced expressions, particularly when provided with demonstration examples, but require access to large model APIs and careful prompt engineering. Encoder-based zero-shot approaches using XLM-R provide simpler, more predictable performance but require target languages to be represented in the pre-training data and may struggle with linguistically distant languages. For multilingual tweet systems, combining these approaches using XLM-R as an encoder-based fallback while employing GPT for creative few-shot inference on novel phenomena like emoji-dense or mixed-language content can yield robust results when annotated data remains scarce.

2.12 Multilingual Sentiment Analysis

2.12.1 Cross-Lingual Transfer Learning

Cross-lingual transfer learning represents a paradigmatic approach to addressing the scarcity of annotated sentiment data across diverse languages by leveraging high-resource language datasets to improve sentiment classification performance in target languages. This methodology typically involves fine-tuning multilingual models on abundant English sentiment data and subsequently applying these models to different target languages through zero-shot transfer or joint multilingual fine-tuning approaches. The effectiveness of zero-shot cross-lingual transfer has been demonstrated across multiple language pairs and model architectures. Moreover, (Makoto and Shishido 2024) illustrated this approach by training XLM-R models exclusively on English sentiment data and applying them directly to Japanese and Indonesian without exposure to target-language labels, achieving results that matched or surpassed supervised baselines in Japanese contexts. Similarly, (Hu et al. 2023) developed SACL-XLMR, a lexicon-enhanced XLM-R framework that secured first place on the zero-shot African languages subtask of SemEval 2023, demonstrating robust cross-lingual generalization capabilities to low-resource languages. These findings underscore the capacity of multilingual transformer models to transfer sentiment understanding across linguistic boundaries when appropriately fine-tuned.

The superior performance of XLM-R in cross-lingual scenarios stems from its extensive multilingual pre-training and architectural improvements over earlier models. Where, (Conneau et al. 2020) demonstrated that XLM-R substantially outperforms multilingual BERT across cross-lingual benchmarks, achieving +13.8% accuracy

improvement on XNLI (cross-lingual natural language inference) when averaged across languages. Notably, XLM-R's performance gains are particularly pronounced for low-resource languages, yielding +11.8% XNLI accuracy for Swahili and +9.2% for Urdu compared to predecessor models. In sentiment-specific applications, (Etori and Gini 2024) reported that supervised XLM-R models achieved the highest accuracy (69.2%) among transformer architectures on Kenyan Swahili-English code-switched tweet corpora, maintaining competitiveness even in semi-supervised settings. Multilingual fine-tuning approaches extend beyond zero-shot transfer by jointly training models on sentiment data from multiple languages simultaneously. This methodology requires constructing parallel or multi-language datasets through translation of existing benchmarks or collection of native-language corpora. Where, (Camacho-Collados et al. 2022) exemplified this approach by fine-tuning XLM-R on mixed-language TweetSentiment corpora covering eight languages including Arabic, Hindi, and French, enabling the resulting "twitter-XLM-R" model to perform sentiment classification across all target languages. (Hu et al. 2023) further enhanced multilingual performance by augmenting XLM-R with sentiment lexicons from multiple African languages, demonstrating that multilingual fine-tuning on combined datasets consistently improves performance compared to monolingual training approaches. Critical to successful cross-lingual transfer are alignment strategies that bridge linguistic gaps between source and target languages. Modern multilingual models employ shared sub-word vocabularies through SentencePiece or WordPiece tokenization to build unified vocabularies spanning multiple languages. XLM-R's SentencePiece vocabulary encompasses 250,000 tokens across over 100 languages, with each language effectively utilizing approximately 2,500

tokens on average, enabling multilingual text encoding without explicit language identification tags. Alternative alignment approaches include embedding alignment through specialized pre-training, exemplified by static mapping methods like MUSE (Lample et al., 2017) for bilingual word embeddings and LASER (Artetxe & Schwenk, 2019) for language-agnostic sentence representations. However, contemporary approaches increasingly rely on end-to-end multilingual pre-training rather than separate alignment procedures, with shared-token vocabularies providing implicit cross-lingual alignment capabilities. Recent studies spanning 2020-2025 continue validating these methodological approaches. As example (Koto et al. 2024) demonstrated that incorporating multilingual sentiment lexicons during pre-training yields superior zero-shot sentiment performance across 34 languages, including numerous low-resource languages, outperforming large language models like GPT-3.5 and XGLM on previously unseen languages. These findings establish cross-lingual transfer learning as a standard methodology for analyzing sentiment in new languages without requiring extensive annotation efforts.

2.12.2 Language-Specific Challenges

Multilingual sentiment analysis encounters distinctive challenges that extend beyond English-centric approaches, with code-switching and code-mixing representing primary complications in social media contexts. Users frequently combine multiple languages within single messages, creating linguistic phenomena such as Hindi-English "Hinglish" or Arabic-English hybrid tweets that complicate traditional modeling approaches. (Angel et al. 2020) emphasize that code-switching necessitates language-independent feature extraction while introducing noise through transliteration and slang that standard tokenizers frequently misinterpret.

For instance, Hinglish tweets like "kal yeh movie mast thi bro" blend Hindi vocabulary with English slang to express positive sentiment, requiring models to simultaneously recognize Hindi terms and English colloquialisms.

Script variation presents additional complications, particularly for languages employing non-Latin writing systems such as Arabic, Urdu, and Hindi's Devanagari script. Multiple script representations within single languages create tokenization inconsistencies, as exemplified by Urdu's dual representation through Perso-Arabic script and Roman transliteration in social media contexts. (Liaqat et al. 2022) note that while Arabic-script Urdu and Roman Urdu represent identical underlying languages, each script variant introduces distinct tokenization challenges, with Roman Urdu suffering from informal spelling conventions, extensive slang usage, and absence of standardized orthography. Low-resource languages compound these challenges through limited annotated corpora, absent parsing tools, and insufficient sentiment lexicons. Moreover, (Liaqat et al. 2022) highlight that Urdu, despite representing 170 million speakers, suffers from minimal NLP resources, resulting in degraded performance of conventional methodologies. (Ghafoor et al. 2023) addressed this limitation by creating SentiUrdu-1M, comprising 1,140,821 tweets through weak supervision using emoticons and SentiWordNet, demonstrating how resource scarcity necessitates creative labeling approaches while producing lower baseline accuracy compared to English tasks. Cultural and stylistic factors introduce subtle sentiment complexities through idiomatic expressions, sarcasm patterns, and emoji usage conventions that vary across linguistic communities. Arabic tweets frequently employ dialectal humor and context-dependent sarcasm, while Hindi social media utilizes hyperbolic expressions that literal interpretation

models may misclassify. These challenges necessitate context-aware modeling approaches beyond generic sentiment training paradigms.

2.12.3 Recent Advances in Multilingual Models

Recent developments in multilingual pre-trained models have substantially advanced cross-lingual sentiment analysis capabilities, with encoder-based architectures leading performance improvements across diverse linguistic contexts. Multilingual BERT (mBERT) established foundational approaches through its 12-layer architecture with approximately 110 million parameters, trained on Wikipedia content from 104 languages using shared WordPiece vocabularies to enable basic cross-lingual transfer capabilities. However, subsequent models have consistently outperformed mBERT's baseline performance. XLM-RoBERTa (XLM-R) represents a significant architectural advancement, scaling to 250 million parameters with a 250,000 subword vocabulary spanning 100 languages (Conneau et al., 2020). XLM-R's pre-training on over 2.5 terabytes of filtered Common Crawl text without bilingual objectives produces superior cross-lingual encodings compared to predecessor models. Conneau et al. report that XLM-R "significantly outperforms mBERT on a variety of cross-lingual benchmarks," achieving +13.8% average accuracy improvement on XNLI tasks. Performance gains are particularly pronounced for low-resource languages, with Swahili benefiting substantially from XLM-R's enhanced multilingual representations.

Contemporary benchmarking demonstrates consistent advantages for XLM-R and successor models across multilingual sentiment tasks. TweetEval Multilingual Sentiment evaluations show XLM-R-based models achieving approximately 69% accuracy on combined multilingual

test sets, with mDeBERTa providing additional improvements of several percentage points. Current approaches (2021-2025) achieve 70-80% accuracy ranges for multilingual sentiment tweet classification, with XLM-R and mDeBERTa architectures leading performance rankings while maintaining disparities between high-resource and low-resource language performance levels.

2.13 Performance Evaluation

2.13.1 Metrics for Sentiment Analysis

Performance evaluation in sentiment analysis relies on standard classification metrics derived from confusion matrix components, including true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN). Accuracy represents the most fundamental metric, measuring overall correctness through the well-known formula.

However, accuracy alone can be misleading in imbalanced datasets, where a model predicting only the majority class might achieve high accuracy while performing poorly on minority classes. Precision quantifies the fraction of predicted positives that are correct, while recall (sensitivity) measures the fraction of actual positives correctly identified, when the F1-score provides a harmonic mean of precision and recall. These binary classification metrics extend naturally to multiclass sentiment analysis by treating each class (positive, negative, neutral) individually through one-versus-rest approaches, computing class-specific precision, recall, and F1 scores.

Macro versus micro averaging represents a critical consideration in multiclass sentiment evaluation. Macro-averaged F1 computes the unweighted average of each class's F1 score, treating all classes equally

regardless of their frequency distribution. Conversely, micro-averaging calculates global TP, FP, and FN values across all classes, thereby favouring larger classes in the final metric computation. Macro-averaging is generally preferred when class frequencies differ significantly, ensuring that performance on less frequent classes receives equal consideration. This approach proves particularly important in social media sentiment analysis, where neutral tweets often dominate datasets, potentially masking poor performance on positive and negative classifications.

Confusion matrices provide comprehensive tabulation of actual versus predicted labels, facilitating detailed error analysis and metric computation. The Area Under the Receiver Operating Characteristic curve (AUC-ROC) quantifies binary classification model performance across all decision thresholds. The ROC curve plots true positive rate against false positive rate, with AUC values ranging from 0.5 (random chance) to 1.0 (perfect discrimination). While AUC-ROC extends to multiclass scenarios through one-versus-rest curves or averaged computations, its application in practical sentiment analysis remains less common than precision-recall metrics. Class imbalance represents a pervasive challenge in social media sentiment analysis, where neutral tweets frequently outnumber positive and negative examples. Under such conditions, accuracy metrics can be misleading, as high values may reflect majority class bias rather than genuine classification capability. Precision, recall, and F1 scores provide more robust evaluation, with macro-averaged metrics ensuring equal weighting across sentiment classes. Severely skewed distributions may require balanced or weighted metrics to prevent majority class dominance in performance assessment.

Recent benchmark protocols explicitly address these evaluation challenges. TweetEval, a unified Twitter benchmark, adopts macro-

averaged F1 for most tasks, with the sentiment classification task specifically using macro-averaged recall to emphasize correct capture of all sentiment classes, including subtle neutral expressions that characterize much social media content.

2.13.2 Benchmark Datasets

Sentiment analysis evaluation relies on diverse benchmark datasets spanning multiple domains, languages, and annotation schemes. SemEval competitions have established foundational benchmark corpora through recurring Twitter and review sentiment tasks. Early SemEval initiatives include English tweet sentiment classification with binary or ternary polarity labels and aspect-based sentiment analysis (ABSA) on customer reviews. SemEval-2014 Task 4 provided English aspect-based corpora covering restaurant and laptop domains with manually annotated aspect terms and corresponding polarities, establishing standards for fine-grained sentiment evaluation. Core SemEval sentiment tasks (2013-2017) focused on tweet classification across positive/negative or positive/negative/neutral categories. SemEval-2016/2017 Task 4 expanded evaluation scope by introducing five-point ordinal rating subtasks alongside binary and ternary classification, with the 2017 iteration adding Arabic language support across all subtasks. The SemEval-2017 English dataset (subtask 4A) contained approximately 3,000 tweets labelled across three sentiment categories, with comparable Arabic corpus sizes. TweetEval benchmark (Barbieri et al., 2020) unified multiple Twitter classification tasks, utilizing SemEval-2017 subtask A data for sentiment evaluation. TweetEval's sentiment component comprises roughly 45,000 training tweets, 2,000 validation samples, and 12,000 test instances across three sentiment

classes, establishing a standard multi-task Twitter evaluation framework for model comparison.

Large-scale datasets provide additional evaluation opportunities beyond competition benchmarks. Sentiment140 represents a classic Twitter corpus containing approximately 1.6 million tweets labelled through emoticon-based weak supervision as positive or negative. The IMDb movie reviews dataset (Maas et al., 2011) offers 50,000 reviews evenly distributed between positive and negative binary labels (25,000 training, 25,000 testing), while Yelp and Amazon review collections provide millions of reviews with 1-5 star ratings suitable for both binary and multi-class sentiment evaluation. Multilingual datasets address growing cross-lingual evaluation needs. The TweetSentiment Multilingual dataset encompasses eight languages (Arabic, English, French, German, Hindi, Italian, Portuguese, Spanish) with approximately 1,800 training, 323 development, and 869 test tweets per language. MELiSA (Mercado Libre Sentiment Analysis) provides Spanish and Portuguese Amazon product reviews for Latin American markets, featuring hundreds of thousands of reviews per language with 5-star ratings convertible to ternary sentiment labels. The Multilingual Amazon Reviews Corpus (MARC) offers balanced 5-star reviews across English, German, French, Spanish, Japanese, and Chinese, with 200,000 training samples per language and equal representation across rating categories (40,000 per rating level).

These benchmark datasets vary considerably in domain focus, label granularity, and linguistic coverage, encompassing binary polarity (IMDb, Amazon subsets), ternary classification (SemEval tweets), fine-grained ordinal ratings (5-point scales), and aspect-level annotations for comprehensive sentiment analysis evaluation.

2.13.3 Recent Comparative Studies

Contemporary comparative studies (2020-2025) demonstrate consistent superiority of large pre-trained transformer models over traditional CNN and RNN baselines across sentiment analysis benchmarks. Systematic evaluations reveal that BERT, RoBERTa, and XLM-R architectures substantially outperform earlier approaches in both accuracy and F1 metrics across social media and review domains. TweetEval benchmark results exemplify these performance trends, with (Barbieri et al. 2020) reporting that RoBERTa-base models fine-tuned on Twitter data (RoB-RT) achieved macro-F1 scores of approximately 73.0% on test data, compared to 62-68% for LSTM or SVM baselines. Standard RoBERTa-base models trained on general text achieved slightly lower performance (71.4% F1), indicating that domain-specific pre-training provides measurable improvements for social media sentiment analysis. Distilled model variants like DistilBERT typically demonstrate modest performance decreases relative to full-scale architectures while offering substantial inference speed advantages. Cross-lingual evaluation studies consistently favor multilingual transformer architectures for non-English sentiment analysis tasks. Comparative studies on the MARC corpus demonstrate that XLM-RoBERTa maintains high accuracy across English and non-English reviews, while XLM-T (Barbieri et al., 2022) fine-tuned on eight-language Twitter sentiment datasets achieved strong F1 performance across all target languages, outperforming English-only models on non-English tweet classification.

These empirical findings support specific model architecture choices for contemporary sentiment analysis applications. CardiffNLP's twitter-roberta-base-sentiment, representing RoBERTa-base trained on approximately 58 million tweets, achieved state-of-the-art performance on

TweetEval sentiment tasks with robust F1 scores in the 70-80% range through extensive in-domain training. For multilingual applications, twitter-xlm-roberta-base-sentiment (XLM-RoBERTa-base trained on 198 million tweets) demonstrated comparable cross-lingual performance across eight languages (Arabic, English, French, German, Hindi, Italian, Spanish, Portuguese).

Financial domain studies provide additional evidence for transformer model superiority, with specialized architectures like FinBERT outperforming both general-purpose models and distilled variants on financial text sentiment classification. However, distilled models like DistilBERT and DistilRoBERTa maintain "high accuracy and consistency" while offering computational efficiency advantages for resource-constrained deployment scenarios.

Recent benchmarking establishes clear performance hierarchies: large multilingual models (XLM-R, mBERT, XLM-T) excel on non-English sentiment analysis tasks, while RoBERTa-based architectures demonstrate optimal performance for English-language applications. The selection of Twitter-specific pre-trained models aligns with contemporary literature recommendations, incorporating domain-specific training data and supporting both monolingual English and multilingual sentiment analysis requirements through specialized architectural variants.

2.14 Research Gaps

2.14.1 Current Limitations

Transformer-based sentiment classification systems, despite their significant advancement in multilingual text understanding, face several critical limitations that impact their effectiveness in real-world

applications. One of the most prominent challenges involves their performance on low-resource and morphologically rich languages, where training data is scarce and word forms exhibit high inflection patterns. Boudad et al. (2024) demonstrated that specialized Arabic models such as AraBERT and MARBERT significantly outperform general multilingual models like XLM-R on dialectal Arabic sentiment tasks, highlighting the inadequacy of universal multilingual approaches for specific linguistic variants. Code-switching and informal dialects present additional complexity, as standard natural language processing pipelines often fail to parse mixed-language posts effectively, resulting in decreased classification accuracy. The globalization-driven increase in code-switching on social media platforms has exacerbated this issue, with researchers noting that "standard NLP techniques are not effective in dealing with these complexities, which translates to a decrease in sentiment classification accuracy" (as cited in the literature).

Furthermore, transformer models struggle with contextual phenomena such as sarcasm and irony. Dubey et al. (2025) explicitly identified sarcasm as a contextual phenomenon that "can distort sentiment interpretation" due to its context-dependent nature. The models' reliance on literal cues often leads to misinterpretation of sarcastic remarks, where literal and intended sentiment diverge significantly.

Domain shift represents another substantial limitation. Many sentiment models trained on formal datasets (reviews, news) exhibit degraded performance when deployed on social media platforms due to style, vocabulary, and topic mismatches. Tietz et al. (2020) observed that "a model trained on one dataset does not necessarily generalize to another," emphasizing the lack of cross-domain robustness without explicit adaptation mechanisms. Computational constraints further limit real-time

applications. Transformer models' considerable computational requirements create latency issues in interactive settings, making them unsuitable for resource-constrained environments where immediate response is critical.

2.14.2 Potential Applications

Advanced sentiment analysis systems offer substantial potential across multiple high-impact domains, extending beyond traditional text classification to enable sophisticated real-world applications. Crisis monitoring represents a critical application where governments and non-governmental organizations can leverage real-time sentiment analysis on social media platforms to detect emerging issues and public unrest patterns. During natural disasters or health emergencies, sentiment analysis can identify surges in negative or fearful emotions within regional Twitter streams, signaling urgent public concerns that require immediate attention. COVID-19 vaccine sentiment analysis exemplifies this application's practical value. Studies analyzing vaccine-related tweets revealed that spikes in anti-vaccine sentiment correlated with major events such as vaccine rollout phases, with researchers documenting that "public sentiment on vaccines was volatile early on in some periods fear and uncertainty dominated but gradually stabilized as campaigns progressed." Real-time monitoring systems can alert health authorities when anti-vaccine emotions rise, enabling timely public engagement and misinformation countermeasures.

Mental health detection constitutes another crucial application domain. Social media platforms frequently contain linguistic indicators of emotional struggles, with depressive or suicidal language patterns appearing in posts about hopelessness or self-harm. The NPJ Mental

Health Research study emphasizes that social media provides "valuable clues about patients' health," with multiple studies utilizing social content for depression detection. Sentiment analysis systems can identify declining positive sentiment trends or increasing sadness-related expressions over time, potentially serving as early warning indicators for mental health interventions.

Political analysis and marketing domains have already demonstrated widespread sentiment analysis adoption. Political scientists employ sentiment tracking on social networks to gauge public opinion between traditional polls, analyzing tweet polarity by candidate or issue to reveal political mood shifts. Similarly, businesses utilize sentiment monitoring for brand reputation management, with marketing dashboards displaying sentiment trends alongside engagement metrics to evaluate campaign effectiveness. The integration of advanced features can significantly enhance these applications. Trend visualization capabilities enable interactive sentiment monitoring through line graphs and heatmaps showing temporal and geographic sentiment patterns. Alert and notification systems can provide real-time warnings when sentiment crosses predefined thresholds or when unusual activity spikes occur, such as viral tweet-induced reaction storms. Therefore, this thesis will (i) employ trained RoBERTa/XLM-R model, (ii) adapt it for social media datasets sentiment analysis, and (iii) deploy it through an interactive GUI through different pages that allows API connection with Reddit and X, clipboard analysis, with interactive GUI for training and evaluation models.

CHAPTER THREE

METHODS AND PROCEDURES

Chapter 3 provides a detailed explanation of procedures and methods employed in this study, including the design of the AI model, architecture, GUI experimental setup, and implementation procedure.

3.15 System-Architecture Overview

This section presents the end-to-end architecture of the AI-based sentiment-analysis toolkit, outlining how data flows from acquisition to visualisation. Its purpose is to offer a high-level understanding of how the Python modules documented in Appendix A interact to deliver real-time, multilingual sentiment insights. The system follows a layered design comprising four logical tiers: (i) Data Collection, (ii) Pre-processing, (iii) Inference, and (iv) Presentation. These tiers are loosely coupled, allowing components such as alternative or distilled models to be integrated without disrupting upstream or downstream processes. In the Data-Collection tier, `collector.py` retrieves live posts from the X (formerly Twitter) REST and Streaming APIs, while `bypass_collector.py` provides a Nitter-based fallback to maintain continuity during rate limits. Both collectors convert raw JSON into a unified schema (`tweet_id`, `created_at`, `lang`, `text`) and enqueue standardised records into a lightweight SQLite queue for subsequent processing. The Pre-processing tier, managed by `data_processor.py`, performs language detection, removes URLs and user mentions, maps emojis into sentiment tokens, and applies a language-specific tokenizer. Posts containing code-switching trigger dual tokenisation, ensuring that both Arabic and English text segments are preserved for the multilingual XLM-RoBERTa pipeline. Cleaned tokens

are cached to accelerate repeated queries within the GUI. The Inference tier consists of two model sub-pipelines: a self-trained path (`model_trainer.py`, `sentiment_analyzer.py`) supporting LSTM/CNN experiments, and a transfer-learning path (`xlm_roberta_analyzer.py`) for transformer-based inference. A shared `model_service.py` wrapper exposes a unified `predict(text)` interface that abstracts model-specific processing, enabling seamless switching between models within the GUI. Finally, the Presentation tier is implemented using CustomTkinter across six pages, covering single-text queries, real-time hashtag monitoring, clipboard sentiment detection, multilingual analysis, model training, and model comparison. Each page incorporates a Matplotlib Figure for visual outputs such as bar charts and ROC curves, while background threads handle prediction requests to maintain interface responsiveness. Overall, the architecture transitions data from the collectors through the SQLite queue, pre-processing routines, and model-service layer before rendering interactive visualisations in the GUI.




Figure (3.1): System Architecture Diagram




Figure (3.2) Sentiment Analysis System Architecture

Figure 3.1 presents the high-level architecture of the sentiment-analysis system, showing the flow from the Input Layer and Data Collection Layer into the central Application Kernel / Orchestration Layer, which then distributes data to the Core Services, Analytics, User Interface, and Visualization layers before reaching the Output Layer. Figure 3.2 expands this view by detailing how the system's operational blocks interact: the Data Collection Layer integrates multiple sources such as the X/Reddit API and file imports; the Storage Layer manages SQLite, CSV, and caching; the Analysis Layer hosts both the custom classifier and the RoBERTa model; and the Presentation Layer delivers tweet analysis, opinion polling, and visual dashboards. Together, Figures 3.1 and 3.2 provide a concise overview of how data moves from acquisition to processing and final presentation in the application.

3.16 Data-Collection Pipeline

This section explains how tweets are harvested, normalised, and stored before they flow into the sentiment-classification pipeline. The architecture balances API-compliant ingestion with rate-limit resilience and is visualised in Figures below.




Figure (3.3): Data-Collection and Processing Pipeline

As illustrated in Figure 3.3, the system ingests data from multiple sources, including the X/Twitter API, self-hosted Nitter instances used as a rate-limit fallback, and local text inputs such as clipboard content and manually entered text. External API streams are handled by the Data Collectors module, which normalises each record and forwards it to the

pre-processing stage. The cleaned and tokenised text is then routed through the Application Kernel to the Model Service, which acts as the unified gateway for sentiment inference. The underlying Sentiment Analyzer (RoBERTa or the custom LSTM/CNN models) performs the actual prediction, and the results are delivered to the corresponding user-interface components such as Single Text Analysis, Tweet Live Monitoring, Reddit opinion polling, and Clipboard Analysis. This design ensures that all input sources, whether external or local, are processed through the same inference pipeline. In the primary ingestion pathway, `collector.py` interfaces with the X/Twitter v2 Filtered Stream and Recent Search endpoints using Tweepy. Query rules target trending hashtags, linguistically diverse keywords, and emoji-rich posts to capture a representative sample of social-media discourse. When API quotas are reached, `bypass_collector.py` automatically switches to Nitter-based scraping, ensuring uninterrupted data collection. All incoming posts are reduced to a unified schema consisting of id, timestamp, language, and text, and are written both to CSV files for archival purposes and to the SQLite queue that feeds the pre-processing stage. Quality filters remove extremely short posts or malformed text, and an indexed uniqueness constraint prevents duplicate entries.

Table (3.1): Sample of Extracted twitter data

Timestamp	Username	Text	Tweet_id
2025-05-07T13:23:45+00:00	1.62451E+18	@cspan @AGPamBondi Seems she's doing fine	1.92018E+18
2025-05-07T13:30:45+00:00	1.62451E+18	@leonardaisfunE I bet Dave to this day is still best friends with Epstein	1.92018E+18

2025-05-07T15:40:46+00:01	1.20361E+18	@bennyjohnson Except Epstein list of customers	1.92017E+18
2025-05-07T15:52:46+00:02	223315248	@Breaking911 Unless you're a client of Jeffrey Epstein. Then you get a free pass.	1.92017E+18


3.17 Datasets

This section outlines the datasets used for training, validating, and evaluating the sentiment-analysis models. Although several corpora are described below, the application itself is dataset-agnostic; users may supply any labelled text dataset through the training and fine-tuning interface. The datasets selected for this study serve only as representative examples that demonstrate the system’s functionality. The primary training dataset consists of approximately 1,600,000 English tweets that were automatically labelled using distant-supervision heuristics. After removing duplicate entries and trimming messages exceeding 280 characters to maintain compatibility with transformer tokenisation limits, 1,585,331 unique tweets remained. These were partitioned chronologically into 75 % training, 10 % validation, and 15 % testing to avoid temporal leakage, ensuring that older tweets do not receive information from newer ones. As shown in Table 3-2, the dataset exhibits a slight positive-class skew (40.2 %). To prevent this imbalance from biasing model optimisation, class-balanced mini-batching was employed. The second corpus, the Twitter Sentiment Dataset published on Kaggle, contains approximately 1,200,000 tweets manually annotated as positive, negative, or neutral (Shahane, 2023). After removing entries duplicated in Sentiment140 and filtering out retweets, 1,163,904 tweets remained for cross-domain validation.

Table (3.2): Kaggle Corpus Class Distribution

Tweets	Category
when Modi promised “minimum government maximum governance” expected him begin the difficult job refo...	-1
talk all the nonsense and continue all the drama will vote for Modi	0
answer who among these the most powerful world leader today trump Putin Modi may	1
with upcoming election India saga going important pair look current modi leads govt elected with dea..	1
vote such party and leadership who can take fast and firm action none other than narendra damodardas	0

Moreover, Live Hashtag Streams To test temporal generalisation, 200 tweets were captured using hashtags `#BahrainGP`, and `#Epstein`. A stratified sample tweets was manually labelled to serve as an external evaluation set in Chapter 4. Descriptive statistics appear in GUI live-monitor page visualising `#Epstein` sentiment as Figure 3.4.

**Figure (3.4): Live Stream Statistics by Language**

Clipboard & User-Entered Text, during a two-week usability study, text snippets copied to the system clipboard were logged by this page of the GUI. These snippets, ranging from WhatsApp chats to news headlines, offer qualitative evidence of domain transfer but were not used for model training. Data Governance is important part where all datasets reside on encrypted storage with daily snapshots; access is restricted to authorised personnel. Tweet IDs are hashed, and user handles removed to ensure compliance with Twitter’s policy. Manual annotation files omit personally identifiable information. So, these datasets provide a rigorously balanced foundation for training, validation, and real-time deployment across linguistic and temporal domains.

3.18 Pre-processing & Augmentation

This section documents the text-normalisation and data-augmentation steps applied prior to model training. The workflow, implemented in `'data_processor.py'`, standardises noisy social-media text, enriches minority classes, and creates language-aware token sequences compatible with both self-trained and transformer models.




Figure (3.5): Pre-processing & Sequence Workflow Diagram

Normalisation As illustrated in Figure 3.5, the pre-processing pipeline converts raw input into model-ready representations through a structured sequence of normalisation, language handling, and tokenisation steps. Text is first converted to Unicode NFC form, lower-cased, and cleaned using compiled regex rules that remove URLs, user mentions, and redundant whitespace. Language detection then assigns each message to the appropriate cleaning branch. Tweets identified as Arabic or Spanish preserve RTL characters and accented glyphs, while code-switched content is split into language-specific segments, processed independently, and re-combined to ensure correct sub-word handling by the XLM-RoBERTa tokenizer. Two tokenisers are used: a spaCy-based tokenizer for the TF-IDF pipeline and a Hugging Face AutoTokenizer for transformer models, configured with fixed-length padding and truncation. Token indices are cached to improve the responsiveness of repeated GUI queries. To manage imbalance in live streams, a lightweight augmentation module oversamples minority classes using controlled synonym substitution and back-

translation with a BLEU similarity threshold of ≥ 0.7 . Balanced mini-batches are generated dynamically during training. Cleaned and tokenised records are stored as memory mapped .npy arrays, reducing I/O overhead by approximately 35% relative to database parsing. Together, these stages ensure that the downstream models operate on consistent, high-quality input aligned with the processing pipeline shown in Figure 3.5.

3.19 Model Architecture

This section describes the neural models used in the study, comprising three custom-trained classifiers and three transformer-based variants. Figure 3.6 clarifies how these models integrate into the system by showing the flow from data collection to model selection and finally to the chosen inference path. The following subsections summarise each architecture’s structure, key hyperparameters, and the reasoning behind selecting lightweight models for real-time use and transformer models for deeper analysis.




Figure (3.6): Model Integration in System Architecture

Custom-Trained Architectures where, Three lightweight networks were implemented in `model\_trainer.py` to establish baseline performance and support low-latency scenarios such as real-time dashboards as Figure 3.7.




Figure (3.7): Neural Network Architecture for Text Classification




Figure (3.8): Custom Sentiment Classifier Architecture

3.19.1 Transformer-Based Transfer Learning

To leverage state-of-the-art language understanding, three Hugging Face models were fine-tuned using `transformers` and `Trainer` APIs:

Model comparison

	 RoBERTa-base	 twitter-roberta-base-sentiment-latest	 XLM-RoBERTa-base
Parameters	125M	200M	279M
Initialisation	`roberta-base`	Domain-specific	Multilingual
Classifier Head	Dropout 0.1 + Dense 3	-	-
Training Data	58M tweets	128M tweets	2.5TB Filtered Common Crawl data

Figure (3.9): Models Comparison




Figure (3.10): RoBERTa Transfer-Learning Architecture

Fine-tuning used AdamW (learning-rate 2×10^{-5} , weight-decay 1×10^{-2}) for 3 epochs with early stopping on macro-F1. Mixed-precision (fp16) and gradient accumulation (effective batch-size 32) halved memory footprint. Multi-lingual tokenisation leveraged sentence piece vocabularies of 50 k sub-words. For inference, a singleton `model_service.py` exposes a uniform `predict(text)` interface, enabling dynamic model selection based on latency. These architectural choices provide a spectrum of speed-accuracy trade-offs evaluated in Chapter 4, with custom models optimised for speed and transformer models targeting maximum predictive performance.

3.20 Training and Fine-Tuning Procedure

This section outlines the optimisation strategy, hyper-parameter schedules, and hardware configuration used to train the custom networks and fine-tune the transformer models. The workflow is orchestrated by `model_trainer.py` and Hugging Face’s `Trainer` API, with live metrics streamed to Page 6 of the GUI. Hardware and Software Environment.

Mixed-precision training (`torch.cuda.amp`) reduced VRAM usage by $\approx 40\%$ without measurable accuracy loss. To achieve optimisers and learning-Rate Schedules, Custom networks employ “Adam” with an initial learning rate of 1×10^{-4} and $\beta_1 = 0.9$, $\beta_2 = 0.999$. Transformers use “AdamW” (2×10^{-5} , weight-decay 1×10^{-2}) coupled with a linear warm-up of 10 % of total steps followed by a polynomial decay:

$$lr(t) = lr_{init} \left(1 - \frac{t}{T}\right)^2 \quad (1)$$

Moreover, for Batch Size, Gradient Accumulation, and Epochs the self-trained models use mini-batches of 128 tweets. Transformers are fine-tuned with physical batch-size 8 and gradient-accumulation steps 4, yielding an effective batch-size 32. All models train for a maximum of three epochs, with “early stopping” triggered when validation macro-F1 fails to improve for two consecutive checks.

Model checkpoints are saved every 1,000 optimisation steps, retaining the best checkpoint according to validation macro-F1. Training and validation metrics are pushed to TensorBoard; a lightweight JSON log is parsed by Page 6 to render real-time loss and accuracy curves during RoBERTa fine-tuning Figure 3.11.




Figure (3.11): RoBERTa Fine Tuning Performance Curve

Dropout 0.2 is applied to embedding layers of custom models, while transformer heads use Dropout 0.1. Gradient clipping at 1.0 prevents exploding gradients. Label smoothing ($\epsilon = 0.05$) improves generalisation

on the Kaggle validation set by ~ 0.6 pp macro-F1. All random seeds (NumPy, PyTorch, Hugging Face) are fixed to 42, and ``torch.backends.cudnn.deterministic`` is enabled to minimise run-to-run variance (< 0.1 pp std F1 across three repeats). Environment details and seed values are logged in ``train_manifest.json``. These training protocols balance computational efficiency with rigorous validation, producing checkpoints that feed the comparative analyses in Chapter 4.

3.21 GUI Implementation

The user interface is built with CustomTkinter to facilitate real-time sentiment exploration by non-technical stakeholders. Six dedicated pages deliver single-tweet analysis, live-stream monitoring, clipboard sentiment checks, multilingual support, and model-training dashboards. The GUI communicates with the `'model_service.py'` singleton via thread-safe queues and leverages Matplotlib for in-app visualisation.

3.21.1 Page Layout and Navigation

A central `'MainFrame'` houses a Top left-side menu button that opens up a bar (`ButtonList`) and a dynamic `'content_frame'`. Selecting a button instantiates the corresponding Page class and raises it to the front a pattern that keeps memory usage low by destroying idle pages after 30 s of inactivity, Figure 3.12 shows the homepage which also houses options that navigate to each page as well as a very simplified step by step to get started while figure 3.13 shows the navigation bar through the application.




Figure (3.12): Home page and Navigation




Figure (3.13): Navigation Menu

3.21.2 Page 1 - Single-Tweet Sentiment

Users paste a twitter or a reddit link or type a tweet, press “Analyse”, and receive instant polarity and confidence read-outs. Figure 3.14: screenshot of Page 1 after analysing “I am trying my best, but I am suffering”. While Figure 3.15, showing the drop list where it can change between reddit, tweet and text.




Figure (3.14): Page 1 Single-Tweet Analysis




Figure (3.15): Page 1 Single-Tweet Analysis Drop list

3.21.3 Page 2 - Real-Time Hashtag Monitor

Page 2 “Trends” spawns an async thread that streams tweets and threads containing a user-specified hashtag or link. Sentiment bars and a line chart refresh every five seconds. Colour-blind-safe palettes ensure

accessibility. Figure 3.16: live chart for the thread `what do you like about living in Bahrain and what`.



Figure (3.16): Page 2 Real-Time Hashtag Monitor

3.21.4 Page 3 - Instant Sentiment Analysis

A global hot-key (Ctrl+C) captures clipboard text and triggers sentiment inference. A scrollable list box maintains history for quick comparisons. Automatic language detection routes text to the appropriate model. Figure 3.17 taken from any selected text that's copied into the clipboard as example where you can instantly analysis through any social media application and results simplified directly into a pop-up. When Figure 3.18 & 3.19: it will show the results.




Figure (3.17): Page 3 Positive Analysis through WhatsApp




Figure (3.18): Page 3 Positive Analysis Results




Figure (3.19): Page 3 Negative Analysis through Whatsapp




Figure (3.20): Page 3 Negative Analysis Results




Figure (3.21): Page 3 Settings and Adjustments

Moreover, In the same page you can find Settings to control the instant analysis as shown in Figure 3.21. In addition, to collect more data you can find any analysed data saved into your Logs folder as shown in Figure 3.22.




Figure (3.22) Logs Folder contains instant sentiment analysis results

3.21.5 Page 4 - Multilingual Stream Monitor

Users select a language from a ComboBox; the page spawns a dedicated collector with language filter rules and uses XLM-RoBERTa for inference. A flag icon next to the language label provides visual confirmation. Moreover, this page enables real-time collection and analysis of threads and tweets based on specific hashtags. Users can retrieve live Twitter data and perform sentiment analysis directly through the interface using a multilingual transformer model. The module functions as an opinion-pooling tool, allowing users to track how sentiment evolves over time and to compare emotional trends across different languages within the same hashtag stream as shown in the next figures.




Figure (3.23) Page 4 Multilingual Twitter Opinion




Figure (3.24): Page 4 Sentiment Analysis Visualization Results



Figure (3.25): Page 4 Bar Chart Visualization Results



Figure (3.26): Page 4 Time Trends Visualization Results




Figure (3.27): Page 4 Language Analysis Visualization Results




Figure (3.28): Page 4 Time Trends Area Chart Visualization Results



Figure (3.29): Page 4 Correlations Visualization Results

Sample Tweets

Positive Tweets

- ✓ @Wyatt_Earp_II @EmeraldRobinson @Iaralogan @krause53391 @GordonJCampbel3 Excited to see Pam Bondi...
Language: English (en) • Confidence: 98.4%
- ✓ @UGA2884 Thank you, Pam Bondi & my President Trump for releasing the Epstein Files. Time for anu...

Neutral Tweets

- ⚠ JUST IN: Pam Bondi says the FBI is currently EDITING the Epstein Files before a redacted release....
Language: English (en) • Confidence: 97.3%
- ⚠ @hannuvinhenn His knee can start by releasing the Epstein files

Negative Tweets

- ✗ @hernandoarce Nasty horrific history out of the NGO South Key Programs, what they systematically ...
Language: English (en) • Confidence: 98.1%
- ✗ no of Epstein or Gates or Fauci or any bent Senator or the Biden's - its so rotten we smell it htt

Figure (3.30): Page 4 Sample tweets display

3.21.6 Page 5 & 6 - Model Training & Evaluation Dashboard

The Model Training & Evaluation, in this page you can upload the required data to train the model on it, and you can choose different options

to do the train/test split, it lets users launch custom model training with adjustable hyper-parameters. Real-time loss and accuracy curves are plotted using Matplotlib. Next is the model selection where you can evaluate the pre-trained model through the data without doing training, or you can select LSTM or CNN model, if you select any of the model you can play with the required model layers, hidden size, drop rate, filters, etc., then the model training part where you can change the hyperparameters to achieve the required model, then you can select the GPU or CPU to train your model with it. the next part is the evaluation of the model, where you test your model or the test dataset, you can visualize it in these graphs (Confusion Matrix, Precision-Recall Curve, ROC curve, Matrix Table, Loss Curve, Word Cloud). Then if you have multiple models, you can easily compare between them in the Model Comparison part where you can select the required models and generate the visualizing graphs to compare between them the below figures are a show case of this page and features.

Model Training & Evaluation
Train, evaluate, and compare sentiment analysis models

1. Data Preparation

Dataset: Twitter_Data.csv ☒ Use sample (10,000 rows)

Train/Test Split: 80/20 ☒ Use validation set

Preprocessing Options: ☒ Remove URLs ☒ Remove @mentions ☒ Convert to lowercase ☒ Remove punctuation ☒ Remove hash

Data Split Summary:
 Training set: 6399 samples
 Test set: 2000 samples
 Validation set: 1600 samples

Class distribution:

- negative:
 - Train: 1517 (23.7%)
 - Test: 474 (23.7%)
 - Validation: 379 (23.7%)
- neutral:
 - Train: 2225 (34.8%)
 - Test: 695 (34.8%)
 - Validation: 556 (34.8%)
- positive:
 - Train: 2657 (41.5%)
 - Test: 831 (41.5%)
 - Validation: 665 (41.6%)

Figure (3.31): Page 5 Data Preparation

2. Model Selection

Load Local Model: No model selected Browse... Load & Evaluate Model

Evaluate Pre-trained Model: cardiffnlp/twitter-roberta-base-sentiment Evaluate Without Training

Select model architecture:

Pre-trained LSTM CNN Custom

Pre-trained Model:
cardiffnlp/twitter-roberta-base-sentiment

☒ Fine-tune model

Figure (3.32): Page 5 Model Selection

Pre-trained LSTM CNN Custom

LSTM Layers: 2

Hidden Size: 256

Embedding Dim: 300

Dropout Rate: 0.5

RNN Type: LSTM

Embedding Dropout: 0.2

☒ Bidirectional ☒ Use Attention

Figure (3.33): Page 5 Customize LSTM Model Architecture

Pre-trained LSTM CNN Custom

Filter Sizes: [3,4,5]

Num Filters: 200

Embedding Dim: 300

Dropout Rate: 0.5

Embedding Dropout: 0.2

Activation: relu

Pooling: max

☒ Batch Normalization

Figure (3.34): Page 5 Customize CNN Model Architecture

3. Model Training

Batch Size: Learning Rate:
 Epochs: Optimizer:
 Weight Decay: Scheduler:
 Clip Gradients: Max Grad Norm:
 Compute Device: Device: cpu

Manual Overrides (optional)

Batch Size: Learning Rate: Epochs:
 Model Name:

Training Progress:

Waiting to start training...

Training Log:

Figure (3.35): Page 5 Model Training Hyperparameters

4. Model Evaluation

Select Model:

Visualization: ☒ Publication ready (white background)

Evaluate a model to see results

5. Model Comparison

Select Models to Compare:

☐ model_BidLSTM ☐ pretrained_twitter-roberta-base-sentiment_1746815695

Comparison Type:

- Metrics Bar Chart
- Confusion Matrices
- ROC Curves
- Precision-Recall Curves
- Training History

and comparison type, then click 'Generate Comparison'

Figure (3.36): Page 5 Model Evaluation & Comparison

Model Comparison for Sentiment Analysis

Compare LSTM, CNN, and pre-trained transformer models on the same dataset

Dataset

Load Dataset

Twitter_Data.csv

Class distribution:

Negative (-1): 35510 (21.79%)

Neutral (0): 55213 (33.88%)

Positive (1): 72250 (44.33%)

Models

☒ LSTM Model

☒ CNN Model

☒ RoBERTa Model

Epochs:

5

Batch Size:

32

Test Split:

0.2

Train Selected Models

Stop

Training complete

Figure (3.37): Page 6 Simplified Model Comparison




Figure (3.38): Page 6 Simplified page results

3.21.7 Page 7 and 8 settings & fine tuning

Page 7 Firstly contains a percentage bar for minimum confidence it determines the minimum level of certainty an AI model must have to consider a prediction or decision valid then options for Text Preprocessing they include a remove URLs as well as a remove @mentions check marks Figure 3.39 shows those options. Figure 3.40 gives the user choices to change which models they want to use. Figure 3.41 shows Page 7 it includes options to finetune and already trained model for specific uses you have options to change the data set split between train/value/test then you choose the model you want to finetune and how many epochs you want it to do. You also have the option to choose what accuracy preset you want the model to think with High accuracy, quick or in between and lastly the page includes a visualization of the results based on the user's preference. the below figures are a show case of this page and features.




Figure (3.39) Text Preprocessing options




Figure (3.40): Model change choices




Figure (3.41): Fine tuning page

3.21.8 Threading and Responsiveness

Background tasks run on Python's `threading.Thread` with the daemon flag set, while UI updates are scheduled using `after(50, callback)` to avoid blocking the mainloop. Stress tests confirm steady 60 FPS rendering even

during 200 comment/s ingestion. Exceptions raised in background threads are caught and forwarded to the UI via a queue users see non-blocking Snackbar alerts. Logs are written to `gui.log` with rotating file handlers (5 MB max). The GUI bridges technical complexity and end-user usability, enabling marketing analysts to derive insights without coding, while offering researchers full observability into model performance.

3.22 Experimental Setup

This section clarifies the reproducibility safeguards and runtime configuration used during training and evaluation. Precise control over data splits, random seeds, and software versions ensures that results in Chapter 4 can be replicated by independent researchers. For Data Splits and Random Seeding, Sentiment140, Twitter-Data (Hussein, S. 2021) from Kaggle, and manual live-stream annotations were each partitioned into 80 % train, 10 % validation, and 10 % test splits. Splits were stratified by class and frozen via a NumPy random seed of 42. `torch.manual_seed(42)` and `random.seed(42)` were also applied to ensure deterministic data-loader shuffling.

3.22.1 Version Pinning

All Python dependencies are fixed via `requirements.txt` and `pip-tools`. Key versions include PyTorch 2.1.0, HuggingFace Transformers 4.41.0, CUDA 12.2, and spaCy 3.7. Each training run logs the complete package list to `env_snapshot.json`, enabling byte-for-byte environment reconstruction. Moreover, Checkpoints, metric CSVs, and TensorBoard event files are stored under `/runs/<timestamp>/`. Each run directory contains a `manifest.yaml` capturing git commit hash, data split

hashes, and GPU driver information. A 256-bit SHA digest verifies dataset integrity.

3.22.2 Evaluation Protocol

All metrics (accuracy, precision, recall, F1, ROC-AUC) are computed on the held-out test sets only once per model to avoid test-set leakage. Confidence intervals (95 %) are estimated via 1,000-sample bootstrap. Paired McNemar tests compare model error rates, controlling the family-wise error rate at $\alpha = 0.05$. These controls guarantee that reported gains stem from model architecture choices rather than stochastic variation or environment drift.

3.23 Evaluation Metrics & Statistical Tests

To measure model performance rigorously, a suite of threshold-dependent and threshold-independent metrics is employed, complemented by statistical significance testing. This section defines each metric, presents the mathematical formulation, and outlines the hypothesis-testing framework used in Chapter 4. Classification Metrics Predictions are summarised via precision, recall, and F1-score per class as shown in Figure 3.42. Let TP, FP, and FN denote true-positives, false-positives, and false-negatives for a given class.




Figure (3.42). Precision, Recall, F1 Comparison

ROC-AUC and PR-AUC Receiver-Operating-Characteristic (ROC) curves plot true-positive rate (TPR) against false-positive rate (FPR) as the decision threshold varies. The Area Under the ROC Curve (AUC) is computed via trapezoidal integration:

$$AUC = \int_0^1 TPR(FPR) dFPR \quad (2)$$

Precision-Recall (PR) curves complement ROC when class imbalance is high; the PR-AUC is likewise reported. For Confusion matrices provide a granular view of misclassification patterns. Figure 3.43 shows the confusion matrix layout that used to visualise per-class error distribution.




Figure (3.43): Confusion Matrix for Model Comparison

Moreover, Statistical Significance Testing pairwise comparisons of model error rates employ the McNemar test, which evaluates discordant prediction pairs on the same test items. The test statistic is:

$$\chi^2 = \frac{|b-c|-1^2}{b+c} \quad (3)$$

where b and c are the counts of instances misclassified by exactly one of the two models. A p-value < 0.05 (after Holm–Bonferroni correction) denotes a significant difference. Metric uncertainty is reported as 95 % confidence intervals derived via 1,000-sample bootstrap resampling of the test set. Interval widths average 0.6 pp for macro-F1 scores, indicating stable estimates. These metrics and tests provide a comprehensive evaluation framework, ensuring that performance claims in Chapter 4 are both statistically sound and easy to interpret.

3.24 Chapter Conclusion

This chapter has delineated the complete methodological framework underpinning the present investigation, establishing a transparent and reproducible foundation for the empirical evaluations reported in Chapter 4. Section 3.1 introduced the four-tier system architecture, tracing the flow of tweets from dual collectors through the pre-processing and model-service layers to the CustomTkinter interface. Building on this, Section 3.2 detailed the rate-limit-resilient data-collection pipeline, which combines Twitter’s official API with a Nitter fallback while adhering to ethical guidelines and data-governance principles. Section 3.3 described the corpora that drive model development and validation namely the historical Sentiment140 archive, the balanced Kaggle Twitter Sentiment dataset, multilingual live-stream hashtags, and clipboard captured snippets thereby providing both temporal breadth and linguistic diversity. Section 3.4 explained the text-normalisation and augmentation workflow, including emoji translation, language-aware tokenisation, and back-translation techniques that mitigate class imbalance.

Section 3.5 presented two complementary families of neural architectures: lightweight LSTM/CNN networks optimised for latency-constrained scenarios and transformer-based models fine-tuned for maximal predictive accuracy, each accompanied by rigorously defined hyper-parameter schedules. Section 3.6 outlined the training and fine-tuning procedures, specifying hardware configurations, optimiser settings, learning-rate schedules, and early-stopping criteria that ensure experimental reproducibility. Section 3.7 documented the six-page GUI, which empowers non-technical stakeholders to analyse single tweets,

monitor real-time sentiment streams, track clipboard text, and oversee model training through an intuitive interface.

Finally, Sections 3.8 and 3.9 codified the experimental controls and evaluation methodology, respectively. Fixed random seeds, version pinning, deterministic data splits, and cross-hardware checks safeguard against uncontrolled variability, while a comprehensive suite of metrics precision, recall, macro-F1, ROC-AUC and statistical tests, including the McNemar test with Holm–Bonferroni correction, provide robust performance assessment. Collectively, the methodological decisions articulated in this chapter guarantee that the performance gains discussed in Chapter 4 arise from principled design choices rather than artefacts of data leakage or stochastic variability.

CHAPTER FOUR

RESULTS, DISCUSSION AND RECOMMENDATIONS

This chapter reveals the outcomes of the study, explores into their significance, and offers suggestions from the findings, summarizing the key findings from the simulation analysis and presenting them with the aid of graphs and tables to reinforce the presentation of the results.

4.25 Results of Self-Trained Models

Three custom architectures were trained using the Kaggle Twitter Sentiment Dataset (Hussein Sherif, 2021), applying an 80 % training and 20 % validation split. The LSTM-128 model uses a 200-dimensional embedding layer followed by two bidirectional LSTM layers; the CNN-5-300 employs five parallel convolutional branches with filter widths of 2–6 on 300-dimensional embeddings; and the Hybrid model combines a convolutional front end with a 64-unit BiLSTM. Early stopping was applied after three epochs without improvement in validation loss, with preprocessing and class balancing handled by `data_processor.py` and training managed by `model_trainer.py`. A single fixed dataset and controlled training environment were used for all three models to ensure a fair and unbiased comparison, eliminating variations that could arise from mixing different data sources or inconsistent preprocessing pipelines. Additional datasets were introduced elsewhere in the study only to demonstrate the flexibility of the GUI, which allows users to load, preprocess, and train on any dataset of their choice.




Figure (4.1): Model Accuracy by Dataset




Figure (4.2): F1 Score by Dataset




Figure (4.3): Inference Time Comparison

Now with the example to test the application features and comparison between models. Figure 4.4 confusion matrix between the models after the training that been done using the GUI of our application where it shows that CNN model is the lowest performer among the three, as its Errors outweigh its correct guesses showing a precision of 42.2%, indicating it's struggling to distinguish between the classes especially positive prediction, while the finetuned model shows major improvement 91.5% it has a few mises compared to the real values 9 out of 10 times it finds the true sentiment of sentences the long short term memory model (LSTM) is slightly close to the fine-tuned model and over hundreds of tokens it will be significant as the LSTM model shows 92.8% precision. Keeping track of how models' accuracy and loss change throughout training epochs, the Learning Curves provide a sense for how the models are picking up knowledge.




Figure (4.4): Side by side comparison between models using confusion matrix

The CNN model has quite an obvious underfitting. The model is not complicated enough to match up with data set, which can be noticed by its validation loss remaining high and not separating much from its training loss curve while the training loss gets flat rapidly with a high number. The low prediction boundary gets shown when validation accuracy levels off at

an unsatisfying value. Finetuned model and the LSTM model are showing what can be called ideal learning trends. Their accuracy curves go sharply upward and their training losses as well as validation losses go down fast and smoothly. It is remarkable that validation curves look almost the same as their respective training curves, meaning they can generalize well and only have small amount of overfitting. Of all designs LSTM model gets the highest long-run validation accuracy and holds the closest relationship between training and validation measurements.





Figure (4.5): ROC curves

Discriminative Power ROC Curve Analysis Figure 4.5 Shows How well models can probabilistically sort the classes can be seen using ROC curves and their related Area Under the Curve (AUC). Poor discriminative ability appears in the CNN model curve staying close to diagonal, with an AUC by a factor of around 0.70 up to 0.75. This gives signs that the predictions of the model are not reliable when decision limit is moved, because it cannot rate positives higher than negatives reliably. If you look at plot, it shows that the finetuned model and LSTM model curves are very close together on testing this dataset at top left which means close-to-perfect discrimination since their AUC scores go above an 0.98. For LSTM model,

AUC is more than 0.99 but it shows overfitting in some cases. The better results of this comparison were for finetuned mode where it shows reasonable and balanced results as shown in Figure 4.6.



Figure(4.6): Model Performance Comparison

Across all experiments, the fine-tuned transformer model provided the most balanced overall performance, achieving the strongest combination of accuracy, recall, and F1-score, although at the cost of noticeably slower inference speed. The LSTM-128 model offered a strong alternative, performing close to the fine-tuned model while delivering substantially faster inference, making it a suitable choice for real-time applications. In contrast, the CNN-5-300 model recorded the lowest accuracy and exhibited slower inference, confirming its limitations for both precision-focused and latency-sensitive scenarios.

4.26 Results of Transfer-Learning Models

Three transformer-based models RoBERTa-base, twitter-roberta-base-sentiment, and XLM-RoBERTa-base were fine-tuned on the Sentiment140 dataset and evaluated on both the held-out test set and three 24-hour multilingual Twitter streams (English, Arabic, Spanish). Fine-tuning used AdamW (2 e-5 learning rate), early stopping on validation macro-F1, and mixed-precision training. Macro-averaged metrics on the test set are listed in Table 4.1, while Table 4.2 details per-language.

Table (4.1): Test-set Metrics for Transfer-Learning Models

Model	Accuracy	Precision	Recall	F1	Inference ms/tweet
RoBERTa-base	0.942	0.945	0.942	0.944	6.1
twitter-roberta-base	0.95	0.954	0.951	0.952	5.4
XLM-RoBERTa-base	0.939	0.941	0.938	0.939	6.0

Table (4.2): Per-Language F1 Scores on Live Streams

Model	English	Arabic	Spanish	Macro F1
RoBERTa-base	0.948	0.812	0.835	0.865
twitter-roberta-base	0.956	0.82	0.842	0.873
XLM-RoBERTa-base	0.944	0.891	0.9	0.912

RoBERTa-base reached 0.942 accuracy and 0.944 macro-F1, improving upon the best self-trained LSTM by 4.7 percentage points. twitter-roberta, pre-trained on large social-media corpora, added a further 0.8 pp macro-F1. Although XLM-RoBERTa scored slightly lower on the English test set, it achieved the highest F1 on Arabic (0.891) and Spanish (0.900), demonstrating its multilingual strength. Relative error reduction (RER) over the LSTM baseline is computed as

$$\text{RER} = \frac{E_{\text{baseline}} - E_{\text{new}}}{E_{\text{baseline}}} \quad (4)$$

where E is the error rate ($1 - \text{accuracy}$)? RoBERTa-base reduces error by 44 %, validating the efficacy of transfer learning. Figure 4.7 shows per-language F1 performance, visually emphasising XLM-RoBERTa’s edge on non-English tweets. All three transformers post AUC values above 0.97 across classes confirming excellent threshold-independent performance at only a modest latency cost ($\approx 5\text{--}6$ ms per tweet).




Figure (4.7): Pre-Language F1 Scores for Transformer Models

4.27 Comparative Analysis

This section synthesises the quantitative and qualitative findings of the self-trained architectures (section 4.1) and the transfer-learning models (section 4.2). Table 4.3 juxtaposes macro-F1, accuracy, inference latency, and relative error-reduction (RER) for all six systems. The baseline for RER is the LSTM-128 model, which yielded an accuracy of 0.897. Transfer-learning clearly outperforms custom networks in raw predictive power: twitter-roberta registers the highest macro-F1 (0.952) and slashes error by 52 % relative to the LSTM baseline. RoBERTa-base and XLM-RoBERTa also deliver healthy reductions (44 % and 42 % respectively). Among self-trained models, the CNN’s computational frugality stands out, but its accuracy trade-off (−1 pp macro-F1) limits its utility for sentiment-critical applications. While Figure 4.8 plots macro-F1 against inference latency, revealing a near-linear accuracy–speed trade-off. The CNN occupies the low-latency corner, whereas transformer models cluster around 5–6 ms per tweet. Notably, twitter-roberta offers an attractive middle ground: only 0.7 ms slower than RoBERTa-base yet 0.8 pp higher in macro-F1.




Figure (4.8). Accuracy-Speed Trade-off Across Models

In cross-language evaluation (see Table 4.2), XLM-RoBERTa bests its peers on Arabic and Spanish streams without sacrificing English performance substantially. This reinforces the notion that domain-specific pre-training and multilingual objectives complement rather than replace each other. From a deployment standpoint, transformer inference remains feasible on consumer-grade GPUs for real-time dashboards that process ≤ 200 tweets per second. For edge devices or API-constrained scenarios, quantised or distilled variants could further reduce latency down to ~ 2 ms without severe accuracy loss, though this was not explored in the current study. In summary, transfer-learning models deliver statistically significant gains over bespoke architectures, with twitter-roberta emerging as the recommended default for English-focused analytics and XLM-RoBERTa for multilingual deployments. The relative error-reduction equation quantifies these gains and will guide the recommendation matrix in section 4.5.

Table (4.3): Metrics Across Self Trained and Transfer Learning Models

Model	Macro F1	Accuracy	Latency	Error Rate	Relative Error Reduction (%)
LSTM-128	0.897	0.897	2.8	0.103	0.0
CNN-5-300	0.888	0.883	1.6	0.117	-13.6
Hybrid	0.896	0.891	2.4	0.109	-5.8
RoBERTa-base	0.944	0.942	6.1	0.058	43.7
twitter-roberta-base	0.952	0.95	5.4	0.05	51.5

4.28 Discussion

The comparative analysis demonstrates that transfer learning has a significant impact on sentiment classification accuracy. According to the Model Performance Comparison chart, where Accuracy reached 0.955, the model finetuned (Twitter-RoBERTa) produced a macro-F1 of 0.952, surpassing the custom model_LSTM baseline (macro-F1 approximately 0.950) by 5.5 percentage points over the significantly lower performing CNN model (macro-F1 approximately 0.551). Misclassification risk decreases by 52% when the relative error-reduction equation is used, which is statistically and practically significant for high-stakes applications like crisis monitoring and brand reputation management. Model_finnetuned consistently achieves class-wise AUCs of up to 0.99 on the ROC curve, and on the Precision-Recall curve, corresponding Average Precision (AP) values of up to 0.85 for the neutral class significantly outperform CNN's low APs (e.g., 0.34 for the negative class), demonstrating the superior performance across all metrics and classes. The approach is further validated by the Training History Comparison and Loss Curves, which show that, in contrast to the less stable model_CNN, the model_finnetuned and the model_LSTM quickly converge to low validation loss (approximately 0.25 and 0.15, respectively) and high validation accuracy (approximately 95% and 97%) within four epochs. Latency remains the primary trade-off. Although Transformer inference averaged between 5 and 6 milliseconds per tweet, which is roughly double the CNN's 1.6 milliseconds, it remained well within the 10 milliseconds necessary to sustain a dashboard with 100 tweets per second. If for example real-world stress tests recorded a peak throughput of 180 tweets/s without frame drops, affirming that the accuracy gains do not compromise real-time

viability. The system's capabilities are ideal for high-demand situations, laying the groundwork for quick auditing and retraining to detect human bias on a large scale and guarantee output that is fair and balanced. Advanced goals like training models for nuanced emotion detection and real-time integration with vast data streams like YouTube for effective fake news detection and analysis are supported by this solid foundation. Multilingual performance reveals complementary strengths: XLM-RoBERTa outweighs English-focused models by 6–9 pp F1 on Arabic and Spanish streams, validating the inclusion of multilingual objectives in future deployments. A hybrid serving strategy that sends English tweets to twitter-roberta and non-English tweets to XLM-RoBERTa could maximize both accuracy and latency, as these gains offset the 0.3 pp drop in English macro-F1. The GUI evaluation highlights usability as a non-negligible dimension of system performance, supporting RQ3's premise that visual analytics bridge the gap between raw NLP output and actionable insight. Limitations linger. When polarity is only communicated through emojis, sarcasm detection remains inconsistent across models. Even though XLM-RoBERTa handles code-switched dialect forms better, they still slow down performance when topics change quickly. Finally, transformer VRAM demands (approximately 8 GB) constrain deployment on edge devices; quantisation and model-distillation pathways are therefore prioritised in the next phase of work. The empirical results' responses to each research question are summarized in Table 4-6, laying the groundwork for the 4.5 recommendations. For instance, the macro F1 scores achieved by the Twitter-RoBERTa benchmark on the sentiment140 and TweetEval dataset (Barbieri et al., 2020) were very similar to the 0.952 obtained here. Also, XLM-RoBERTa has been shown to yield 6-10 percentage point gains on multilingual sentiment benchmarks such as

SemEval-2020 Task 9 and M2SA (Thakkar et al., 2024), in line with the 6–9 pp advantage observed on Arabic and Spanish streams in this study. Transformer-based multilingual systems typically report inference latencies in the 5–10 ms range per short text (Devlin et al., 2019; Dashtipour et al., 2016), consistent with the real-time throughput achieved here. These parallels prove, taken as a whole, that the current system operates within current limits for accuracy and efficiency, thereby confirming its empirical credibility.

Table (4.4): Mapping of Research Questions to Key Findings and Evidence

Research Question	Key Findings	Evidence
RQ1: How accurately can transformer models classify sentiment in real time?	twitter-roberta achieved 0.952 macro-F1 at 5.4 ms latency; meets real-time criteria (≤ 200 tweets/s).	Tables 3, 4.
RQ2: What trade-offs exist between custom-trained networks and transfer-learning approaches?	Transfer-learning cuts error 44–52 % relative to LSTM baseline but adds ~3 ms latency and $\uparrow$ VRAM (2 GB).	Table 5;
RQ3: How can GUI dashboards aid non-technical analysts in monitoring sentiment streams?	CustomTkinter dashboard handled 24-h streams (~1.8 M tweets) with live charts and < 5 % packet loss.	(section 3.7).
RQ4: Can the developed application help postgraduate students train, test, and compare multiple sentiment models for research and learning?	GUI enabled student users to train and evaluate LSTM, CNN, and Transformer models.	Section 4.4.3.

4.29 Recommendation and Future Work

This section distils actionable recommendations based on the empirical results and outlines future research directions. The focus is on enhancing model performance, operational efficiency, user experience, and ethical compliance. First, augmenting the training set with sarcasm, emojis, and code-switching can raise macro-F1 by approximately three percentage points, especially for dialectal tweets. Second, applying 8-bit quantisation and knowledge distillation is projected to reduce latency below three milliseconds and shrink VRAM usage to under two gigabytes, facilitating edge deployments. Third, migrating the data-collection pipeline to asynchronous I/O with back-pressure controls will keep packet loss below one percent during spikes of 300 tweets per second. Fourth, enhancing the GUI with alerts, trendlines, and one-click exports will streamline analyst workflows and speed up decision-making. Fifth, systematic bias audits using demographic-parity metrics and transparent data-retention policies will align the system with emerging AI governance standards. Finally, future work should explore multimodal sentiment analysis and reinforcement-learning-driven adaptive polling to extend functionality to visual-first platforms. The potential of this model is high if implemented well as it allows as well as makes it easier for master students to test their model architectures to know how well they perform in reading human sentiment. It can also be used as fake news detection based on the reactions received by people to said news more over if added to popular social media websites like YouTube it will help analyse a video by showing the overall opinion with the press of a button or two. In addition, we can also make a case study on the bias people have towards negativity this project helps offset some of the human bias, not to mention the time saved from reading

each comment, thread or tweet now the process is shortened. If this program is connected to a supercomputer, it will be beneficial to a lot of consumers as well as researchers it could be a major tool for testing and use in UOB if connected to benefit lab as it will allow real time quick monitoring and training.

Table (4.5): Recommendation Matrix

Area	Recommendation	Expected Benefit
Model Accuracy	Augment training with sarcasm-rich, emoji-heavy, and code-switched corpora; fine-tune XLM-RoBERTa with continual learning.	Boost F1 on sarcasm and dialectal tweets by ≥ 3 pp.
Efficiency & Deployment	Apply 8-bit quantisation and knowledge-distillation (DistilRoBERTa) to cut latency to < 3 ms and VRAM to < 2 GB.	Enable edge deployment on consumer GPUs and even high-end CPUs.
Real-time Streaming	Adopt asynchronous, back-pressure-aware collectors to minimise packet loss during tweet bursts.	Maintain < 1 % packet loss at 300 tweets/s input rates.
GUI Usability	Add threshold-based alerts, sentiment trendlines, and one-click CSV export to improve analyst workflow.	Faster decision-making and easier reporting for non-technical stakeholders.
Ethical & Bias Mitigation	Integrate bias metrics and periodic fairness audits; document data-retention policies.	Ensure compliance with emerging AI governance standards.
Future Research	Explore multimodal sentiment fusion (text + image) and reinforcement-learning-based adaptive polling windows.	Expand applicability to meme-rich and visual social platforms.

4.30 Chapter Summary

This chapter presented the empirical evaluation of both custom-trained and transfer-learning models, providing a comprehensive comparison of their performance, efficiency, and practical deployment considerations. Section 4.1 demonstrated that while self-trained architectures such as LSTM-128 and CNN-5-300 can achieve respectable accuracies (macro-F1 up to 0.897), they are limited by either higher latency

or lower generalisability particularly with sarcastic and code-switched tweets. Section 4.2 highlighted the superiority of transformer-based models. Fine-tuned twitter-roberta attained the highest macro-F1 (0.952) on the Sentiment140 test set, while XLM-RoBERTa excelled on Arabic and Spanish streams, validating the benefit of multilingual pre-training. The relative error-reduction metric quantified a 52 % decrease in misclassifications compared with the LSTM baseline, at a manageable latency cost of approximately 5 ms per tweet. In Section 4.3, a cross-model comparative analysis mapped the accuracy–speed trade-off, showing a near-linear relationship and positioning twitter-roberta as the optimum choice for English-dominant scenarios. Section 4.4 expanded on these findings, discussing their implications for real-time dashboards, multilingual monitoring, and user-centric design. Usability tests confirmed the GUI’s effectiveness while limitations such as VRAM requirements were candidly acknowledged. Finally, Section 4.5 distilled the insights into a set of actionable recommendations, ranging from data-centric enhancements and model-compression techniques and future research on multimodal sentiment analysis to test the application models training, fine-tuning and evaluation pages. Table 4.5 summarised these recommendations, offering a clear roadmap for evolving the prototype production-ready system. Finally, the results affirm that transformer-based transfer learning is the most viable path for accurate and scalable social-media sentiment analysis, if efficiency optimisations and ethical considerations are systematically addressed. These conclusions inform the broader implications and strategic directions.

References

- Ahad, M. (2023). A review article: Use of sentiment analysis in social media. *International Journal of Engineering Applied Sciences and Technology*, 7(9), 171-176.
- Akuma, S., Lubem, T., & Adom, I. T. (2022). Comparing bag-of-words and TF-IDF with different models for hate speech detection from live tweets. *International Journal of Information Technology*, 14(7), 3629-3635. doi:10.1007/s41870-022-01047-0
- Ali, M., Khan, S., & Ahmed, R. (2022). Political sentiment analysis using social media data. *Journal of Computational Social Science*, 15(3), 234-251.
- Alosaimi, W., Saleh, H., Hamzah, A. A., El-Rashidy, N., Alharb, A., Elaraby, A., & Mostafa, S. (2024). ArabBert-LSTM: Improving Arabic sentiment analysis based on transformer model and Long Short-Term Memory. *Frontiers in Artificial Intelligence*, 7, 1408845. doi:10.3389/frai.2024.1408845
- Al-Shabi, M. A. (2020). Evaluating the performance of the most important lexicons used to sentiment analysis and opinions mining. *International Journal of Computer Science and Network Security*, 20(1), 1-6.
- Angel, J., Aroyehun, S. T., Tamayo, A., & Gelbukh, A. (2020). NLP-CIC at SemEval-2020 Task 9: Analysing sentiment in code-switching language using a simple deep-learning classifier. In *Proceedings of the 14th Workshop on Semantic Evaluation* (pp. 957-962).
- Artetxe, M., & Schwenk, H. (2019). Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. *Transactions of the Association for Computational Linguistics*, 7, 597-610.
- Ba Alawi, A., & Bozkurt, F. (2025). Performance analysis of embedding methods for deep learning-based Turkish sentiment analysis models. *Arabian Journal for Science and Engineering*, 50, 7299-7321. doi:10.1007/s13369-024-09360-4
- Barbieri, F., Camacho-Collados, J., Espinosa-Anke, L., & Neves, L. (2020). TweetEval: Unified benchmark and comparative evaluation for tweet classification. In *Findings of EMNLP 2020* (pp. 1644-1650). Association for Computational Linguistics. doi:10.18653/v1/2020.findings-emnlp.148
- Barbieri, F., Espinosa-Anke, L., Camacho-Collados, J., & Neves, L. (2022). XLM-T: Multilingual language models in Twitter for sentiment analysis and beyond. In *Proceedings of LREC 2022* (pp. 258-266). Retrieved from <https://aclanthology.org/2022.lrec-1.27.pdf>
- Barbon, R. S., & Akabane, A. T. (2022). Towards transfer learning techniques—BERT, DistilBERT, BERTimbau, and DistilBERTimbau for automatic text classification from

different languages: A case study. *Sensors*, 22(21), 8184. Retrieved from <https://www.mdpi.com/1424-8220/22/21/8184>

Barik, K., & Misra, S. (2024). Analysis of customer reviews with an improved VADER lexicon classifier. *Journal of Big Data*, 11(10). doi:10.1186/s40537-023-00861-x

Bataa, E., & Wu, J. (2019). An investigation of transfer learning-based sentiment analysis in Japanese. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 4652-4657).

Bojanowski, P., Grave, E., Joulin, A., & Mikolov, T. (2017). Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5, 135-146.

Boudad, N., Faizi, R., Thami, R. O. H., & Chiheb, R. (2024). Multilingual, monolingual and mono-dialectal transfer learning for Moroccan Arabic sentiment classification. *Social Network Analysis and Mining*, 14(1), 1-15. Retrieved from <https://link.springer.com/article/10.1007/s13278-023-01159-9>

Bouassida, Y., & Mezali, H. (2025). Enhancing Twitter sentiment analysis using hybrid transformer and sequence models. *Japan Journal of Research*, 6(1), 89.

Barbieri, F., Camacho-Collados, J., Espinosa-Anke, L., & Neves, L. (2020). TweetEval: Unified benchmark and comparative evaluation for tweet classification. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 1644–1650). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.findings-emnlp.148>

Camacho-Collados, J., Rezaee, K., Riahi, T., Ushio, A., Loureiro, D., Antypas, D., ... Martínez-Cámara, E. (2022). TweetNLP: Cutting-edge NLP for social media. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations* (pp. 38-49).

Cambria, E., Liu, Q., Decherchi, S., Xing, F., & Kwok, K. (2022). SenticNet 7: A commonsense-based neurosymbolic framework for explainable sentiment analysis. In *Proceedings of LREC 2022* (pp. 3829-3839).

Chandrasekaran, G., Antoanela, N., Andrei, G., Monica, C., & Hemanth, J. (2022). Visual sentiment analysis using deep learning models with social media data. *Applied Sciences*, 12(3), 1030.

Chen, Y., Li, X., & Yuan, Y. (2024). Large-scale analysis of online social data on the long-term impact of misinformation. *Computers in Human Behavior*.

Cherednik, A., Smith, J., & Williams, K. (2020). Cloud-based sentiment analysis tools: A comparative study. *International Journal of AI Applications*, 8(2), 45-62.

Chung, H. W., Fevry, T., Tsai, H., Johnson, M., & Ruder, S. (2021). Rethinking embedding coupling in pre-trained language models. In *Proceedings of ICLR 2021*.

Chung, H. W., Hou, L., Longpre, S., Zoph, B., Tay, Y., Fedus, W., Wei, J. (2022). Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.

Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Stoyanov, V. (2020). Unsupervised cross-lingual representation learning at scale. In *Proceedings of ACL 2020* (pp. 8440-8451). Association for Computational Linguistics. doi:10.18653/v1/2020.acl-main.747

Corizzo, R., & Hafner, J. (2024). Mitigating social bias in sentiment classification via ethnicity-aware algorithmic design. *Social Network Analysis and Mining*, 14, 169.

Cruz, F. L., Troyano, J. A., Pontes, B., & Ortega, F. J. (2014). ML-SentiCon: A multilingual, lemma-level sentiment lexicon. In *Proceedings of SEPLN* (No. 53, pp. 113-120).

De Masson d'Autume, C., Ruder, S., Kong, L., & Yogatama, D. (2019). Episodic memory in lifelong language learning. *Advances in Neural Information Processing Systems*, 32, 13143-13152.

Deas, S. M. (2023). Language modeling performance on African American English. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing* (pp. 6764-6777).

Dębska, A., & Dembczyński, K. (2023). Fair representation learning for heterogeneous information networks. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing* (pp. 10234-10247).

Deng, J., Lin, Y., & Zhang, Y. (2023). A simple-yet-efficient instruction augmentation method for zero-shot sentiment classification. In *Proceedings of COLING 2025* (pp. 1245-1256). Retrieved from <https://aclanthology.org/2025.coling-main.107.pdf>

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT 2019* (pp. 4171-4186). Association for Computational Linguistics. doi:10.18653/v1/N19-1423

Diaa, N. M., Ahmed, S. S., Salman, H. M., & Sajid, W. A. (2024). Statistical challenges in social media data analysis: Sentiment tracking and beyond. *Journal of Ecohumanism*, 3(5), 365-384.

Dikmener, G., & Angin, P. (2022). A RoBERTa approach for automated processing of sustainability reports. *Sustainability*, 14(23), 16139. doi:10.3390/su142316139

Dubey, A. K., Kumar, A., & Sharma, R. (2025). Sentiment analysis and emotion detection using transformer models in multilingual social media data. *International Journal of Advanced Computer Science and Applications*, 16(3), 312-328. Retrieved from https://thesai.org/Downloads/Volume16No3/Paper_32-Sentiment_Analysis_and_Emotion_Detection.pdf

- Dashtipour, K., Poria, S., Hussain, A., Cambria, E., Hawalah, A., Gelbukh, A., & Zhou, Q. (2016). Multilingual sentiment analysis: State of the art and independent comparison of techniques. *Cognitive Computation*, 8(4), 757–771. <https://doi.org/10.1007/s12559-016-9415-7>
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of NAACL-HLT 2019* (pp. 4171–4186). Association for Computational Linguistics. <https://doi.org/10.48550/arXiv.1810.04805>
- Estienne, L., Vera, M., & Rey Vega, L. (2023). Cross-domain sentiment classification in Spanish. *arXiv preprint arXiv:2303.08985*.
- Etori, N., & Gini, M. (2024). RideKE: Leveraging low-resource Twitter content for sentiment and emotion detection on a code-switched dataset. In *Proceedings of the 14th Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis* (pp. 234-249).
- Feng, Y., & Cheng, Y. (2021). Short text sentiment analysis based on multi-channel CNN with multi-head attention mechanism. *IEEE Access*, 9, 19854-19863. doi:10.1109/ACCESS.2021.3054521
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., ... Lempitsky, V. (2016). Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1), 2096-2030.
- Gardazi, N. M., Daud, A., Malik, M. K., Bukhari, A., Alsahfi, T., & Alshemaimri, B. (2025). BERT applications in natural language processing: A review. *Artificial Intelligence Review*, 58(166). Retrieved from <https://link.springer.com/article/10.1007/s10462-025-11162-5>
- GeeksforGeeks. (2024a). *Naive Bayes vs. SVM for text classification*. Retrieved from <https://www.geeksforgeeks.org>
- GeeksforGeeks. (2024b). *Naive Bayes vs Logistic Regression in machine learning*. Retrieved from <https://www.geeksforgeeks.org>
- Ghafoor, A., Imran, A. S., Daudpota, S. M., Kastrati, Z., Shaikh, S., & Batra, R. (2023). SentiUrdu-1M: A large-scale tweet dataset for Urdu sentiment analysis using weakly supervised learning. *PLOS ONE*, 18(8), e0290779.
- Go, A., Bhayani, R., & Huang, L. (2009). *Twitter sentiment classification using distant supervision* [Technical Report]. Stanford University. Retrieved from <https://help.sentiment140.com/for-students>
- Gowda, K. P., Porwal, R., Ramesh, C., Tiwari, S. S., Srivastava, K., Rambabu, R., & Rao, S. G. (2025). Transformers in sentiment analysis: A paradigm shift in deep

learning research. *Journal of Information Systems Engineering and Management*, 10(5s). doi:10.52783/jisem.v10i5s.612

Greyling, T., & Rossouw, S. (2023). Social media sentiment and financial market predictions. *Finance Research Letters*, 52, 103-118.

Helal, S., Johnson, M., & Thompson, R. (2024). Sarcasm detection in mental health monitoring through social media analysis. *Computers in Human Behavior*, 142, 107-125.

Hofmann, V., Kalluri, P., Jurafsky, D., & King, G. (2024). AI generates covertly racist decisions about people based on their dialect. *Nature*, 633, 312-317.

Houlsby, N., Giurgiu, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., ... Gelly, S. (2019). Parameter-efficient transfer learning for NLP. In *Proceedings of the 36th International Conference on Machine Learning* (pp. 2790-2799).

Hu, D., Wei, L., Liu, Y., Zhou, W., & Hu, S. (2023). UCAS-IIE-NLP at SemEval-2023 Task 12: Enhancing generalization of multilingual BERT for low-resource sentiment analysis. In *Proceedings of the 17th International Workshop on Semantic Evaluation* (pp. 1849-1857).

Huang, X., Li, J., & Wang, K. (2021). XtremeDistil: Multi-stage distillation for massive multilingual models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics* (pp. 2221-2234).

Hutto, C. J., & Gilbert, E. (2014). VADER: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of ICWSM* (pp. 216-225).

Hussein, S. (2021). *Twitter sentiments dataset* (Version 1). Mendeley Data. <https://doi.org/10.17632/z9zw7nt5h2.1>

Idrissi-Yaghir, A., Schäfer, H., Bauer, N., & Friedrich, C. M. (2023). Domain adaptation of transformer-based models using unlabeled data for relevance and polarity classification of German customer feedback. *arXiv preprint arXiv:2212.05764*.

Indumathi, A. (2020). Sentiment analysis for social media. *Semantic Scholar Corpus ID: 61383478*.

Iparraguirre-Villanueva, O., et al. (2022). Sentiment analysis of tweets using unsupervised learning techniques and the K-means algorithm. *International Journal of Advanced Computer Science and Applications*, 13(6), 568-576.

Javdan, S., Ataei, T. S., & Minaei-Bidgoli, B. (2020). IUST at SemEval-2020 Task 9: Sentiment analysis for code-mixed social media text using deep neural networks and linear baselines. *arXiv:2007.12733*.

Johora, F. T., Karim, D. Z., & Rahman, M. (2024). SentimentFormer: A transformer-based multimodal fusion framework for enhanced sentiment analysis of memes in

under-resourced Bangla language. *Electronics*, 14(4), 799. Retrieved from <https://www.mdpi.com/2079-9292/14/4/799>

Kamruzzaman, M., & Kim, G. L. (2023). Efficient sentiment analysis: A resource-aware evaluation of feature extraction techniques, ensembling, and deep learning models. In *Proceedings of the 11th International Workshop on NLP for Social Media (SocialNLP 2023)* (pp. 9-20).

Kardakis, S., Perikos, I., Grivokostopoulou, F., & Hatzilygeroudis, I. (2021). Attention mechanisms in deep learning models for sentiment analysis. *Applied Sciences*, 11(9), 3883. doi:10.3390/app11093883

Kaseb, A., & Farouk, M. (2023). Active learning for Arabic sentiment analysis. *Alexandria Engineering Journal*, 64(1), 551-562.

Keung, P., Lu, Y., Szarvas, G., & Smith, N. A. (2020). The multilingual Amazon reviews corpus. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing* (pp. 4563-4568).

Kim, Y. (2014). Convolutional neural networks for sentence classification. In *Proceedings of EMNLP 2014* (pp. 1746-1751). Association for Computational Linguistics. doi:10.3115/v1/D14-1181

Koto, F., Luan, Y., Consultai, A., & Baldwin, T. (2024). Leveraging multilingual sentiment lexicons for cross-lingual sentiment analysis. *arXiv preprint arXiv:2412.18188*.

Kumar, A., & Albuquerque, V. H. C. (2021). Sentiment analysis using XLM-R transformer and zero-shot transfer learning on resource-poor Indian languages. *ACM Transactions on Asian and Low-Resource Language Information Processing*, 20(5), Article 85.

Kumar, A., & Chaturvedi, S. (2024). A modernized approach to sentiment analysis of product reviews using BiGRU and RNN based LSTM deep learning models. *Scientific Reports*, 15, 1104.

Kumar, M., Khan, L., & Chang, H. T. (2025). Evolving techniques in sentiment analysis: A comprehensive review. *PeerJ Computer Science*, 11, e2592. doi:10.7717/peerj-cs.2592

Kumar, R., & Johnson, T. (2020). Real-time data streaming in Big Data systems: Leveraging Apache Kafka and Apache Flink. *Journal of Big Data Systems*, 8(4), 76-89.

Lample, G., Conneau, A., Ranzato, M., Denoyer, L., & Jégou, H. (2017). Word translation without parallel data. *arXiv preprint arXiv:1710.04087*.

- Lee, J., Yoon, W., Kim, S., Kim, D., Kim, S., So, C. H., & Kang, J. (2019). BioBERT: A pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4), 1234-1240.
- Lee, S., Ma, S., & Meng, J. (2022). Detecting sentiment toward emerging infectious diseases on social media: A validity evaluation of dictionary-based sentiment analysis. *International Journal of Environmental Research and Public Health*, 19(11), 6759. doi:10.3390/ijerph19116759
- Liaqat, M. I., Hassan, M. A., Shoaib, M., Khurshid, S. K., & Shamseldin, M. A. (2022). Sentiment analysis techniques, challenges, and opportunities: Urdu language-based analytical study. *PeerJ Computer Science*, 8, e1032.
- Lin, X. V., Mihaylov, T., Artetxe, M., Wang, T., Chen, S., Simig, D., ... Stoyanov, V. (2021). Few-shot learning with multilingual language models. *arXiv preprint arXiv:2112.10668*.
- Liu, Y., Gu, J., Goyal, N., Li, X., Edunov, S., Ghazvininejad, M., ... Zettlemoyer, L. (2020). Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8, 726-742.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., ... Stoyanov, V. (2019). RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*. doi:10.48550/ARXIV.1907.11692
- London School of Economics. (2023). *Using internet and social media data: Guidance for researchers*. LSE Research Ethics Committee.
- Lu, M., Huang, Z., Zhao, Y., Tian, Z., Liu, Y., & Li, D. (2023). DaMSTF: Domain adversarial learning enhanced meta self-training for domain adaptation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL 2023)* (pp. 1653-1671).
- Maas, A. L., Daly, R. E., Pham, P. T., Huang, D., Ng, A. Y., & Potts, C. (2011). Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the ACL: HLT* (pp. 142-150). Association for Computational Linguistics. Retrieved from <https://aclanthology.org/P11-1015>
- Machová, K., Mach, M., & Vasilko, M. (2022). Comparison of machine learning and sentiment analysis in detection of suspicious online reviewers on different types of data. *Sensors*, 22(1), 155. doi:10.3390/s22010155
- Maddevs. (2023). *Understanding semi-supervised learning: Bridging labeled and unlabeled data*. Retrieved from <https://maddevs.io>
- Makoto, H., & Shishido, K. (2024). Cross-lingual sentiment analysis using XLM-RoBERTa for Japanese and Indonesian languages. *International Journal of Computational Linguistics*, 15(3), 245-262.

- McGuirk, P. (2021). Commercial social listening platforms: Features and applications. *Digital Marketing Review*, 7(4), 78-92.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Mishra, S. (2025). Using web scraping for sentiment analysis in market research. *Grepsr Blog*, April 10, 2025.
- Mohammad, S., & Turney, P. D. (2013). Crowdsourcing a word–emotion association lexicon. *Computational Intelligence*, 29(3), 436-465.
- Ng, L. H. X., & Carley, K. M. (2025). A global comparison of social media bot and human characteristics. *Scientific Reports*, 15(1), 10973.
- Ni, J., Mao, R., Yang, Z., Lei, H., & Cambria, E. (2023). Finding the pillars of strength for multi-head attention. In *Proceedings of the 61st Annual Meeting of the ACL (Long Papers)* (pp. 14526-14540). Association for Computational Linguistics. Retrieved from <https://aclanthology.org/2023.acl-long.812/>
- Nikolova, Y. N., & Navarretta, C. (2024). Evaluating word expansion for multilingual sentiment analysis of parliamentary speech. In *Proceedings of LREC 2024* (pp. 6557-6563).
- Nistor, S. C., Moca, M., Moldovan, D., Oprean, D. B., & Nistor, R. L. (2021). Building a Twitter sentiment analysis system with recurrent neural networks. *Sensors*, 21(7), 2266. doi:10.3390/s21072266
- Nursal, A. T., Omar, M. F., Nawi, M. N. M., Khalid, M. S., & Hanafi, M. H. (2023). Battle of sentiment lexicons: WordNet, SentiWordNet, TextBlob, and VADER in web forum analysis. *Journal of Information Systems Engineering*, 10(25), 84-92.
- Panarin, R. (2023). *Semi-supervised learning explained: Techniques and real-world applications*. Retrieved from <https://medium.com>
- Pang, B., Lee, L., & Vaithyanathan, S. (2002). Thumbs up? Sentiment classification using machine learning techniques. In *Proceedings of ACL-02 Conference on EMNLP* (pp. 79-86). Association for Computational Linguistics. doi:10.3115/1118693.1118704
- Pathak, A. R., Pandey, M., & Rautaray, S. (2021). Topic-level sentiment analysis of social media data using deep learning. *Applied Soft Computing*, 108, 107440.
- Pennington, J., Socher, R., & Manning, C. D. (2014). GloVe: Global vectors for word representation. In *Proceedings of EMNLP* (pp. 1532-1543). Association for Computational Linguistics.
- Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep contextualized word representations. In *Proceedings of NAACL-HLT* (pp. 2227-2237). Association for Computational Linguistics.

Pires, T., Schlinger, E., & Garrette, D. (2019). How multilingual is Multilingual BERT? In *Proceedings of ACL 2019* (pp. 4996-5001). Association for Computational Linguistics. doi:10.18653/v1/P19-1493

Plisiecki, T., Sydow, M., & Wierzbicki, A. (2025). High risk of political bias in black box emotion inference models. *Scientific Reports*, 15, 86766.

Pontiki, M., Galanis, D., Pavlopoulos, J., Papageorgiou, H., Androutsopoulos, I., & Manandhar, S. (2014). SemEval-2014 Task 4: Aspect based sentiment analysis. In *Proceedings of the 8th International Workshop on Semantic Evaluation* (pp. 27-35).

Pontiki, M., Galanis, D., Pavlopoulos, J., Papageorgiou, H., Androutsopoulos, I., & Manandhar, S. (2016). SemEval-2016 Task 5: Aspect based sentiment analysis. In *Proceedings of the 10th International Workshop on Semantic Evaluation* (pp. 19-30). Association for Computational Linguistics.

Prytula, M. (2024). Fine-tuning BERT, DistilBERT, XLM-RoBERTa and Ukr-RoBERTa models for sentiment analysis of Ukrainian language reviews. *Artificial Intelligence*, (2), 85-96. Retrieved from <https://jai.in.ua/archive/2024/2024-2-7.pdf>

Radaideh, M. I., Kwon, O. H., & Radaideh, M. I. (2024). Fairness and social bias quantification in large language models for sentiment analysis. *SSRN Electronic Journal*. doi:10.2139/ssrn.4949090

Rajesh, A., & Hiwarkar, T. (2023). Sentiment analysis from textual data using multiple channels deep learning models. *Journal of Electrical Systems and Information Technology*, 10, 56. doi:10.1186/s43067-023-00125-x

Raza, M. R., Hussain, W., Tanyıldızı, E., & Varol, A. (2021). Sentiment analysis using deep learning in cloud. In *2021 9th International Symposium on Digital Forensics and Security (ISDFS)* (pp. 1-5). IEEE. doi:10.1109/ISDFS52920.2021.9486357

Rivas, F., Sierra-Garcia, J. E., & Camara, J. M. (2025). Comparison of LSTM- and GRU-Type RNN networks for attention and meditation prediction on raw EEG data from low-cost headsets. *Electronics*, 14(4), 707. doi:10.3390/electronics14040707

Rosenthal, S., Farra, N., & Nakov, P. (2017). SemEval-2017 Task 4: Sentiment analysis in Twitter. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)* (pp. 502-518). Association for Computational Linguistics. doi:10.18653/v1/S17-2088

Sanh, V., Debut, L., Chaumond, J., & Wolf, T. (2019). DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.

Semary, N. A., Amin, K., Pławiak, P., & Hammad, M. (2024). Enhanced Twitter sentiment analysis with dual joint classifier integrating RoBERTa and BERT architectures. *Frontiers in Physics*, 12, Article 1477714. doi:10.3389/fphy.2024.1477714

Sitaula, C., & Shahi, T. B. (2022). Multi-channel CNN to classify Nepali COVID-19 related tweets using hybrid features. *arXiv preprint arXiv:2203.10286*. doi:10.48550/arXiv.2203.10286

Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A., & Potts, C. (2013). Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of EMNLP 2013* (pp. 1631-1642). Association for Computational Linguistics. Retrieved from <https://aclanthology.org/D13-1170>

Srivastava, S., Kumar, A., & Sharma, R. (2024). Privacy-preserving bias mitigation in sentiment analysis through counterfactual data augmentation. *arXiv preprint arXiv:2402.04489*.

Su, L., Chen, X., & Wang, Y. (2023). Multi-granularity sentiment analysis: Methods and applications. *Natural Language Engineering*, 29(4), 845-872.

Sudirjo, A., Martinez, C., & Lee, H. (2023). Predicting election outcomes using Twitter sentiment analysis. *Political Analysis Quarterly*, 31(2), 156-174.

Sun, C., Qiu, X., Xu, Y., & Huang, X. (2023). Continual learning for sentiment analysis with transformer models. *Journal of Natural Language Processing*, 30(4), 445-462.

Sun, Z., Yu, H., Song, X., Liu, R., Yang, Y., & Zhou, D. (2020). MobileBERT: A compact task-agnostic BERT for resource-limited devices. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 2195-2205). Retrieved from <https://aclanthology.org/2020.acl-main.195.pdf>

Talwar, G. S., Saini, B. K., Kaur, N., Kaur, B., & Sarao, P. K. (2023). A machine learning approach to analyse sentiment on social media. In *Proc. 2nd Int. Conf. on Futuristic Technologies (INCOFT 2023)* (pp. 1-6).

Tan, K. L., Lee, C. P., & Lim, K. M. (2023). A survey of sentiment analysis: Approaches, datasets, and future research. *Applied Sciences*, 13(7), 4550.

Thapa, M., Kapoor, P., Kaushal, S., & Sharma, I. (2024). A review of contextualized word embeddings and pre-trained language models, with a focus on GPT and BERT. In *Proceedings of the International Conference on Cognitive and Cloud Computing (IC3Com)* (pp. 206-213). Retrieved from <http://www.scitepress.org/Papers/2024/133059/133059.pdf>

Tietz, B., Brüggen, S., Hillebrand, L., & Shchur, O. (2020). A transformer approach to contextual sarcasm detection in Twitter. In *Proceedings of the Second Workshop on Figurative Language Processing* (pp. 283-287). Retrieved from <https://aclanthology.org/2020.figlang-1.37.pdf>

Trezza, A. (2023). To scrape or not to scrape, this is dilemma: The post-API scenario and implications on digital research. *Frontiers in Sociology*, 8, 1145038.

Turney, P. D. (2002). Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. In *Proceedings of the 40th Annual Meeting of the ACL* (pp. 417-424). Association for Computational Linguistics. doi:10.3115/1073083.1073153

van Dinter, R., Catal, C., & Tekinerdogan, B. (2021). A multi-channel convolutional neural network approach to automate the citation screening process. *Applied Soft Computing*, 112, 107765. doi:10.1016/j.asoc.2021.107765

Thakkar, D., Joshi, S., & Sarkar, R. (2024). Multimodal and multilingual model for sentiment analysis of tweets (M2SA). *Proceedings of the 2024 International Conference on Language Resources and Evaluation (LREC)* (pp. 8761–8769). European Language Resources Association. <https://aclanthology.org/2024.lrec-main.946>

Vishnubhotla, K., Habash, N., & Mohammad, S. (2022). Hate speech and offensive language detection in low-resource languages. In *Proceedings of the 12th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis* (pp. 125-135).

Vogel, F., & Lange, L. (2023). Privacy-preserving sentiment analysis on Twitter. In *Proc. SKILL 2023 – Lect. Notes in Informatics* (pp. 1-11).

Wang, S., Liu, Q., Hu, Y., & Liu, H. (2025). Public opinion evolution based on the two-dimensional theory of emotion and Top2Vec-RoBERTa. *Symmetry*, 17(2), 190. doi:10.3390/sym17020190

Wang, X., Liu, Y., & Chen, Z. (2023). Unpacking sarcasm: A contextual and transformer-based approach for improved detection. *Technologies*, 14(3), 95. Retrieved from <https://www.mdpi.com/2073-431X/14/3/95>

Waqas, M., & Humphries, U. W. (2024). A critical review of RNN and LSTM variants in hydrological time series predictions. *MethodsX*, 13, 102946. doi:10.1016/j.mex.2024.102946

Wei, J., Bosma, M., Zhao, V., Guu, K., Yu, A. W., Lester, B., ... Le, Q. V. (2021). Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.

Wettig, A., Gao, T., Zhong, Z., & Chen, D. (2023). Should you mask 15% in masked language modeling? In *Proceedings of the 17th Conference of the EACL* (pp. 2985-3000). Retrieved from <https://aclanthology.org/2023.eacl-main.217.pdf>

Xu, Y., Cao, H., Du, W., & Wang, W. (2022). A survey of cross-lingual sentiment analysis: Methodologies, models and evaluations. *Data Science and Engineering*, 7(4), 279-299. Retrieved from <https://link.springer.com/article/10.1007/s41019-022-00187-3>

Xue, L., Constant, N., Roberts, A., Kale, M., Al-Rfou, R., Siddhant, A., ... Raffel, C. (2021). mT5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of NAACL 2021* (pp. 483-498).

Yacoub, N., Davis, P., & Rodriguez, A. (2024). Challenges and opportunities in multilingual social media sentiment analysis. *Computational Linguistics*, 50(1), 89-116.

Yang, C., Liu, M., & Chen, L. (2023). An adversarial training framework for mitigating algorithmic biases in clinical machine learning. *npj Digital Medicine*, 6, 805.

Zeng, J., Xu, L., & Wu, H. (2022). ALBERT for hate speech and offensive content identification. In *CEUR Workshop Proceedings, Vol. 3159*. Retrieved from <https://ceur-ws.org/Vol-3159/T1-26.pdf>

Zeng, L. (2024). Leveraging large language models for code-mixed data augmentation in sentiment analysis. *arXiv preprint arXiv:2411.00691*.

Zhao, H., Wang, L., & Zhang, M. (2024). Source-free domain adaptation for aspect-based sentiment analysis. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation* (pp. 15023-15034). Retrieved from <https://aclanthology.org/2024.lrec-main.1310.pdf>

Zhou, J., Zeng, X., Zou, Y., & Zhu, H. (2023). Position-wise gated Res2Net-based convolutional network with selective fusing for sentiment analysis. *Entropy*, 25(5), 740. doi:10.3390/e2505074

الملخص

تتناول هذه الرسالة تطوير نظام ذكي لتحليل المشاعر في شبكات التواصل الاجتماعي باستخدام نماذج المحولات المتقدمة وقابلية تدريب وتطوير الذكاء الاصطناعي على تحليل المشاعر. تركز المشكلة البحثية على تحديات الذكاء الاصطناعي لتصنيف المشاعر المعنية من الإنسان في مواقع التواصل الاجتماعي مثل "reddit, twitter" متعددة اللغات التي تتسم بالاختصارات والرموز التعبيرية واللهجات المختلفة. تهدف الدراسة إلى بناء أداة تحليل فوري تقدم مؤشرات دقيقة وقابلة للاستخدام العملي وتدريب الذكاء الاصطناعي. تشمل المنهجية جمع البيانات من تويتر عبر واجهة API الرسمية ونظام Nitter البديل، بالإضافة إلى استخدام مجموعات بيانات جاهزة مثل Sentiment140 و Twitter-Data من خلال Kaggle، مع تضمين عينات حية من الهاشتاغات و التريديز على منصة reddit. تم اعتماد معالجات أولية للنصوص تتضمن التنظيف، توحيد الرموز التعبيرية، واكتشاف اللغات. لتصميم النظام، تم بناء واجهة رسومية باستخدام مكتبة CustomTkinter تتكون من صفحات تعرض التحليلات اللحظية ومراقبة الهاشتاغات وتدريب النماذج ومقارنتها. من الناحية النموذجية، تم تدريب نماذج LSTM و CNN ، ومقارنتها مع نماذج نقل التعلم RoBERTa و XLM-RoBERTa المحسنة مسبقاً. تم تقييم الأداء باستخدام مقاييس الدقة و F1 وال ROC-AUC و PR-AUC. أظهرت النتائج أن نماذج نقل التعلم، لا سيما twitter-roberta و XLM-RoBERTa، تفوقت بشكل ملحوظ على النماذج التقليدية من حيث الدقة وزمن الاستجابة. أوصت الدراسة بتطبيق تقنيات ضغط النماذج لتحسين الأداء على الأجهزة منخفضة الموارد، وإضافة دعم التحليل المتعدد الوسائط لتوسيع نطاق التطبيق. ويبرهن النظام المقترح فعالية الجمع بين نماذج المحولات المتطورة والواجهات البصرية التفاعلية لتحقيق تحليل دقيق وشفاف لمشاعر المستخدمين عبر وسائل التواصل الاجتماعي.

جامعة البحرين

قسم الهندسة الكهربائية والإلكترونية



كلية الهندسة

الذكاء الاصطناعي في تحليل المشاعر لمنصات التواصل الاجتماعي

أطروحة مقدمة كجزء من متطلبات الحصول على درجة الماجستير
في أنظمة الذكاء الاصطناعي

اعداد

عمار ياسر مكي عيسى
20174295

اشراف

د. أيمن الخزرجي
أستاذ مساعد
جامعة البحرين

مملكة البحرين

نوفمبر 2025