WORLD COMET RESEARCH GROUP

International Journal of Applied Data Science & Modern Computing
Vol. 5, No. 1, 2022
Doi: 10.67228/30713498/IJADSMC-2022PI5R3M
PP. 01-13

Original Article

Real-Time Sentiment Monitoring Using Social Media Streams

Dr. Nthabiseng Mokoena

Department of Psychology, Pretoria Central University, South Africa.

Received: 02-02-2022 Revised: 02-22-2022 Accepted: 03-01-2022 Published: 03-03-2022

ABSTRACT

The viral increase in social media services has altered the way people, organisations and the government convey their views and respond to events in the real world. These websites are producing huge quantities of unstructured text data on demand, providing opportunities that sentiment analysis and opinion mining have never had. The objective of real-time sentiment monitoring is to provide continuous analysis of such high flowing data streams to detect the emotions, opinions as well as trend among the populace in real time. The competencies are essential in areas like brand reputation management, political analysis, crisis management, financial market prediction and public health surveillance. Nonetheless, sentiment monitoring in real-time causes tremendous technical difficulties. The data streams of social media have high volume, velocity, and variety, as well as noises, linguistic informality, multilingual messages, sarcasm, and concept drift. Such dynamic environments cannot be addressed using traditional batch-oriented sentiment analysis algorithms because they have drawbacks of latency and scalability. As a result, current systems progressively depend on distributed stream processing systems, deep learning, and adaptive learning plans to attain low-latency high-accuracy sentiment inference. The following paper provides an in-depth research of social media stream-based real-time sentiment monitoring. It critically examines the history of sentiment analysis algorithms, starting with lexicon-based and classical machine learning proxies to the state of the art deep learning and transformer-based algorithms. The paper also presents a structure of modular real-time sentiment monitoring framework that incorporates data ingestion, preprocessing, features representation, sentiment classification and visualization layers. Sentiment scoring and streaming classification mathematical formulations are discussed to give a theoretical ground to these. Big-data social media datasets are evaluated through experimental evaluation on the framework of streaming architecture, which proves the suitability of deep learning models in trade-off accuracy and latency. Comparative analysis helps in bringing out the trade-offs of the traditional and neural solution in real-time constraints. The results affirm that properly constructed real-time sentiment monitoring systems can deliver timely actionable information at reasonable computer processing cost. Given that the future prospects in the research focus on multilingual sentiment analysis, explainability and ethical standards of large-scale social media monitoring, a conclusion of the paper is provided on the future research directions.

KEYWORDS

Real-time sentiment analysis, social media streams, opinion mining, deep learning, stream processing, natural language processing.

ISSN: 3071-3498
PP: 01-13.

1. INTRODUCTION

1.1. Background

The fastest spreading of social media services, Twitter, Facebook, and Reddit, has had a paradigm shift in how opinions, feelings, and attitudes are launched, shared, and outsourced in the society. Millions of users constantly produce immense amounts of posts, comments, likes, and reactions consistently that represent the popular opinion on the products, services, political choices, social movements, and world happenings. It is a continuous and live stream of user-generated content, which is the potential source of a rich and vibrant data on the collective behavior analysis and the impact of the time on the development of the collective opinion. Sentiment analysis or opinion mining is a computational problem regarding the extraction of subjective information, sentiment polarity and emotional tone, among text data. The earliest work done in this area was based mainly on thematic datasets and offline analysis where some inference was made after the fact that the data was collected. Yet, the increasing demands of instant and operational intelligence have led to the paradigm shift to the real-time sentiment monitoring systems that have the ability to process the social media streams as they are received. These kinds of systems allow identifying the emerging trends in time, quickly identifying the issues of the population, and predicting possible emergencies. As an example, organizations can gauge the customer sentiment on live product releases, brands can gauge the change in reputation in reaction to a marketing campaign, and the government can gauge the reaction of the population in case of an emergency or a policy announcement. These wide-ranging applications underscore the extreme significance of scalable, low-latency, and precise sentiment analysis frameworks which are capable of being functional in real-time, sentiment monitoring being a fundamental aspect of contemporary data-driven choices.

1.2. Importance of Real-Time Sentiment Monitoring

In the time of constant online communication, real-time sentiment surveillance has become a highly essential feature. Compared to the old-fashioned offline sentiment analysis method, real-time systems can give results about how the opinions of the people change in real-time, thus helping people make more informed decisions on various fields. Real-time sentiment monitoring is important because it can be observed in terms of the following important aspects.

Fig 1 - Importance of Real-Time Sentiment Monitoring

1.2.1. Timely Decision Making

The main benefits of sentiment monitoring in real-time include facilitation of timely and proactive decisions. When conducting analysis of social media streams as they appear, organizations will notice the new trends or abrupt change in the mass opinion, or adverse response on a particular

ISSN: 3071-3498
PP: 01-13.

event very fast. This immediacy ensures that the stakeholders react before problems blow out of proportion and mitigate any possible reputational, financial, or social effects. In high-speed settings, reactions may end up fruitless when the sentiment insight is delayed, thus the need of real-time analysis.

1.2.2. Crisis Detection and Management

Sentiment monitoring in real-time is an important element of detecting a crisis early on as well as managing the crisis. Blips of negative those affect a negative sentiment are usually the initial warning signs that there is a crisis, including a crash in services, product malfunction, dissatisfaction among the population, or the dissemination of false information. The real-time monitoring of sentiments helps organizations and authorities to identify such anomalies early enough and take corrective measures. Such capability is especially useful when it comes to such spheres as public safety, healthcare, or governance, where response time can alleviate the impact of a massive effects.

1.2.3. Business Intelligence and Customer Engagement

In the case of businesses, sentiment analysis is a real time actionable intelligence regarding the perception, preferences and expectations of customers. Sentiment tracking when launching a product, in a promotion campaign, or a service update is one possible way of organizations determining the perception of the people in real time and responding accordingly. It also allows individually engaging customers because it is possible to detect dissatisfaction in real time and to respond timely, enhancing customer satisfaction and brand loyalty.

1.2.4. Societal and Policy Insights

In addition to business use, real-time sentiment tracking can be used to provide useful information on the dynamics of the society and discourse. Governments, policymakers, and researchers are able to gauge the reaction of the people to policies, elections, or emergencies thereby aiding in evidence-based decision making. The real-time sentiment tracking of social events helps in strengthening the perception of a collective mind as the story unfolds and its ability to respond in real-time to events within a society, which is why the use of sentiment monitoring is an essential tool in the digital age.

1.3. Challenges in Social Media Stream Analysis

Although it has great potential, real-time sentiment mining in the stream of social media posts poses several technical and linguistic challenges making the scalable analysis of the matter more complicated. The incredibly high rate of data creation and volume on the social media platforms, where millions of messages can be created in a rather short period is one of the leading challenges. The high rate and volume of continuous data streams of this scale demand a stream processing design that is highly scalable and fault-tolerant, including with low latency to provide an acceptable level of throughput. Real-time systems can face the issue of performance loss and slow insights unless successfully parallelized and managing resources. The other big challenge has been created by the informal and noisy language of the social media. The content produced by users often includes slang, abbreviations, misspellings, emojis, hashtags and non-standard grammar, with the effect that sending sentiment signals becomes more difficult and the standard methods of natural language processing may be less effective. Although emojis and hashtag usually have a high level of emotional indication, their meaning is very much situational and it is hard to extract a similar sentiment across all cases. Also, sarcasm and irony are still a thorn in the flesh, its expressions can become completely

ISSN: 3071-3498
PP: 01-13.

sentimental to the opposite of the meaning of words. This shortcoming may result in misclassification with or without advanced deep learning models. The number of multilingual content makes the social media stream analysis even harder, especially in international media when users are many, and some communicate in various languages and dialects even in the same message. It is not a trivial task to create models that can be effective in cross-linguistic generalization and at the same time with the ability to capture cultural and contextual subtleties. In addition, the sentiment of the social media is extremely dynamic and a problem of concept drift occurs, in which the sentiment phrase, vocabulary as well as the behaviour of the user change with time. Historical data based on which models are trained can become obsolete, and it is necessary to constantly adapt and retrain models. Taken together, it can be concluded that the analysis of social media sentiment in real-time is a complex issue, and a strong, flexible, and scalable analytical structure is necessary.

2. LITERATURE SURVEY

2.1. Lexicon-Based Sentiment Analysis

Lexicon-based sentiment analysis is one of the most intuitively defined ways of opinion mining, based on the use of sentiment lexicons, collections of words or phrases that have been assigned a polarity rating of positive, negative, or neutral. These methods generally calculate document or sentence-level sentiment as an aggregate of the individual token sentiment scores (the weights being frequency or syntactic significance). Due to their lack of a labelled training data, lexicon-based methods lend themselves to the low-resource setting well as to early-stage systems in which annotated corpora are not available. They are appealing in their application because they are transparent and interpretable, which is useful in systems that need to be explained, like decision support systems. Nevertheless, these methodologies are limited to a significant degree because of their simplicity and their high efficiency in terms of computation. They also have difficulty in capturing the contextual features of negation, sarcasm, irony, and the intensity modifiers which may significantly change the sentiment polarity. Moreover, systems based on lexicon are often strongly domain on top of which feelings are carried by words and lexicons do not change their meaning as users evolve language, slang and abbreviations, which seem to be eminent features of social media streams making such systems less robust and particularly more inaccurate in a real-world implementation.

2.2. Machine Learning-Based Approaches

Sentiment analysis as an application of machine learning was an important advancement as compared to lexicon-based methods because, through the use of supervised learning algorithms, sentiment patterns are inferred directly through data. Naïve Bayes, Support Vector Machines and logic regression are the classical examples of classifiers which were successfully adopted because they have a strong theoretical basis, and their computational complexity level is rather low. These models are based on artificial feature representations, such as bag-of-words, n-grams, part-of-speech tags and term frequency-inverse document frequency (TF-IDF), to reflect lexical and statistical properties of text. Machine learning approaches are better adaptive and classification accurate where enough labeled data are present in comparison with lexicon-based methods. Nevertheless, they are highly sensitive to feature quality and domain optimization, and typically they demand large amounts of preprocessing and human feature hand-crafting. Also, the majority of the classical models of machine learning are trained in batch, which restricts their ability to perform in real-time or streaming models where the distribution of sentiment can change over time as a result of concept

ISSN: 3071-3498
PP: 01-13.

drift. The high training frequency and model performance often result in significant overhead on computations and thus make scaling and real-time reaction difficult in large social media analytics.

2.3. Deep Learning and Neural Models

The use of deep learning has led to a tremendous change in sentiment analysis since it allows automatic representations learning and enhances situational comprehension. Convolutional Neural Networks (CNNs) make effective n-gram features that model local features, whereas Recurrent Neural Networks (RNNs) model sequential dependencies and long-range context in text, especially the Long Short-Term Memory (LSTM) networks. These models do not require any massive manual feature engineering, they have always performed better in benchmark datasets compared to the traditional machine learning methods. Most recently transformer-based architectures like BERT have established new performance benchmarks through the use of self-attention mechanisms and contextual word embeddings which dynamically learn the meaning of words given an adjacent context. These models are very successful to addressing language complexities like polysemy and feelings changes with context. Their size, however, and high amounts of computation pose a challenge as to inference timeliness, memory usage and energy cost, especially when running in real time sentiment detection systems. Consequently, the application of deep neural models in low-resource and streaming systems or methods frequently requires compression, distillation, or hardware optimization methods.

2.4. Stream Processing Frameworks

In order to enable real-time longevity sentiment analysis in terms of high velocity data streams, current systems are becoming more compelled to incorporate distributed stream processing designs offering scalability, fault resilience, and low-latency calculation. Apache Kafka is a system that allows reliable ingress and buffering of large volumes of social media data, and the processing engines, such as Apache Spark, can provide real-time analytics using the capabilities of in-memory computing and parallel execution. Such frameworks enable the sentiment models to process continuously on streaming data as opposed to processing static datasets, allowing studies to obtain timely results of prevailing trends and events in the opinion of the population. Stream processing architectures have window-based analysis, micro-batching, and machine learning pipeline integration, which is why they are easily applied to large-scale sentiment monitoring workflows. Nevertheless, it is still difficult to guarantee end-to-end low latency and high accuracy, especially when some complicated deep learning models are used. In turn, as per the recent studies, the balanced consideration of throughput, scalability, and analytical accuracy can be achieved by co-designing efficient sentiment models and stream processing infrastructures.

3. METHODOLOGY

3.1. System Architecture Overview

The real-time sentiment monitoring framework proposed is created as a layered structure to maintain scalability, modularity as well as low-latency processing of textual streams of high velocity. All the layers have a distinct role, which allows integrating data ingestion, analysis, and visualization segments effortlessly and providing real-time sentiment intelligence.

ISSN: 3071-3498
PP: 01-13.

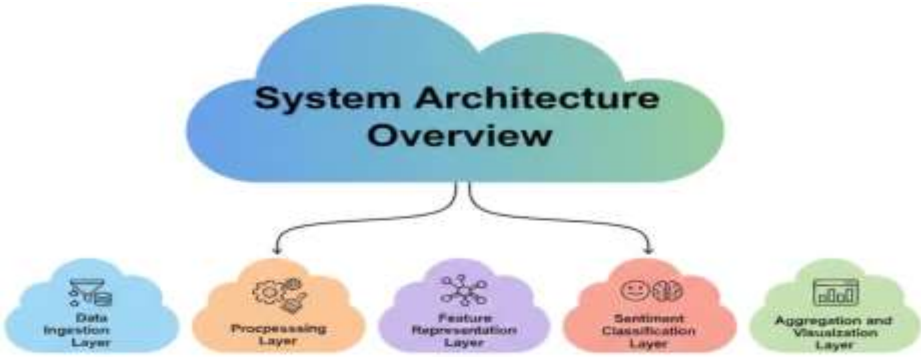

Fig 2 - System Architecture Overview

3.1.1. Data Ingestion Layer

The data ingestion layer is charged with the responsibility of incorporating in real time textual information in high volume that comes in heterogeneous form including social media platforms, online forums, and news feeds. This layer guarantees high availability (via buffering of incoming streams) as well as scalable and fault-tolerant data capture (through handling of the fluctuations in throughput). It can help maintain a continuous data flow and retain the order of messages and reduce the likelihood of losing messages, which is essential in real-time sentiment tracking systems. The ingestion layer decouples data producers and downstream processing elements and allows the system to support sudden increases in data volume without any impairment in performance.

3.1.2. Preprocessing Layer

The preprocessing layer converts the unclean and raw textual data to a clean and normalized one which can be used in analytical modeling. This level entails the process of tokenization, lowercasing, stop word removal, noise filtering, and emojis, hashtag, and URLs that are common in social media text. Besides, the language layer is used to resolve linguistic issues, such as spelling variants, shortenings, and negations, to enhance the accuracy of sentiment signals. Preprocessing should be efficient to eliminate sparsity of the data and provide high quality inputs to the downstream models especially in real-time streaming setup.

3.1.3. Feature Representation Layer

The feature representation layer transforms text that has been preprocessed into the stack of numbers which can be fed into machine learning or deep learning models. Based on the modeling strategy chosen, this layer can either use more classical modeling tools like TF -IDF and n-grams or more modern distributed modeling tools like word embeddings and contextual embeddings. The representation used will have a great effect on the accuracy of classification and computation efficiency. In the real-time systems, this layer is streamlined so as to strike a balance between the expressive power and low latency computation, making it possible to transform streaming text into meaningful feature spaces quite quickly.

3.1.4. Sentiment Classification Layer

The sentiment classification layer is the main layer of analytical power of the framework which the extracted features are each assigned to a sentiment category like positive, negative, or neutral. This layer combines trained sentiment models, which can be classical machine learning classifiers, deep neural and transformer based models. The models are made to develop semantic and context patterns in the written text and are characterized by rapid inferences that can be deployed in real-time. To facilitate adaptability, this layer can adopt periodic updates to the model or online learning systems to deal with changing sentiment pattern.

3.1.5. Aggregation and Visualization Layer

The visualization layer allows identification of high-level sentiment information based on aggregating classification outputs over time windows, topic, or geographical area. It calculates such aggregate measures as sentiment distribution, polarity trends, and anomaly detection indicators to aid in decision-making. These insights are displayed using visualization dashboards that are composed of easy to understand charts and real-time graphs, and by this means the stakeholders can efficiently track the dynamics of public opinion. This layer then converts raw sentiment prediction into actionable intelligence to complete the entire end-to-end pipeline of real-time sentiment monitoring system.

3.2. Data Collection and Streaming

The initial phase of the suggested real-time sentiment monitoring system is data collection and streaming that allows acquiring and supplying social media content in high velocity continuously and processes it for downstream analysis. Data used in social media are largely based on official programs Application Programming Interfaces (APIs) that are programmatic access to user-created content including posts, comments, and metadata. These APIs serve the keyword-based queries, hashtags, and topic filters which enable the system to selectively receive relevant streams as per the objectives of the application. Each message as it is received gets a specific timestamp at the moment of ingestion, indicating the time of event creation or the time of entering the system. Temporal consistency and precise trend analysis is therefore achieved. Timestamping is a key element in ensuring the sequencing of events and it can be used to perform window-based analytics to aggregate sentiments in real-time. After gathering, the data are fed into a distributed streaming pipeline which is set up to process large scale, continuous data streams with low latency. The incoming messages are segmented according to the attributes like the message keys, topics or source identifiers to make several compute nodes work in parallel. Workload distribution and system scaling Partitioning provides a balanced distribution of workload by enabling the framework to handle large throughput streams. Besides, the streaming pipeline also supports the buffering and queueing to absorb the bursty traffic patterns typically witnessed during trending events or breaking news. With fault-tolerance, message replications, and checkpointing, fault-tolerance features serve to avert the loss of data and data processing reliability during the occurrence of node failures. The streaming infrastructure also allows the realization of real-time data enrichment by the use of auxiliary metadata such as source identifiers, language tags and geolocation information whenever available. This augmented stream of data is subsequently streamed to other layers of processing in order to do preprocessing and sentiment analysis. Comprehensively, the data collection and streaming segment provides scalable, resilient and smooth data flow, which is a strategic enabler of real-time sentiment tracking in the unstatic and vast social media settings.

3.3. Text Preprocessing

Real time sentiment monitoring comprises of text preprocessing which is a crucial step since raw social media information tends to be extremely noisy, unstructured and informal. The main goal of preprocessing is to convert the received streams of text to a clean and normal form and one that is free of noise and can be classified properly by taking into account the sentiment carrying information.

ISSN: 3071-3498
PP: 01-13.

The initial step of a preprocessing pipeline is tokenization, which splits the text into separate tokens, e.g. words, symbols, or emojis. Since the language in social media is informal, the use of sophisticated means of tokenization is introduced to properly process contractions, repetitive characters, and a combination of alphanumeric phrases. This part defines the fundamental basis of the further linguistic analysis and subsequent features extracting. Normalization is the next step of tokenization, which tries to minimize variability in lexicon by transforming text into a uniform form. Some of the common normalization methods are lowercasing, extraneous punctuations and normalizing the long words (e.g., eliminating soooo to so). Users URLs, user mentions as well as special symbols are removed or substituted with fixed tokens to stop bringing extraneous noise. The stop-word removal is then used to remove words that recur most and which are weak in terms of their semantics like articles and prepositions which do not normally contribute to the sentiment polarity. This step helps to improve computational efficiency and model attention to sentiment-relevant terms by dimensionality reduction of the text representation.

3.4. Sentiment Classification Model

The sentiment classification part will be assigned the role of establishing an emotional polarity of each social media incoming message based on the semantic content. An ordered sequence of messages is then converted into a given message as a sequence of tokens (each of which represents a word, symbol, or emoji retrieved at the preprocessing stage). It is a contextual representation of the message in the sequence and therefore the dependencies and sentiment signals that might be created by the arrangement of words and by word order and structure of composition are captured by the model. The individual tokens in the sequence get converted into a numerical vector with the assistance of an adequate embedding methodology that can encompass fixed word embeddings or context-sensitive depictions provided by the neural language models. A neural network functionality is then used to process the embedded token sequence that learns to automatically extract sentiment-relevant features. This functionality is characterized with the help of a set of trainable weights, which are optimized at the learning stage with the help of labeled sentiment data. Based on the selected architecture, the neural network can contain convolutional layers to detect patterns where information is localized, recurrent layers to find sequence dependencies, or attention-based layers to attend to tokens pertinent to sentiment. The result of this neural work is a score vector, with each value being the confidence of this model with respect to a certain sentiment category. A softmax operation is used to convert these scores into decipherable predictions. The softmax is used to normalize the output scores into a probability distribution over all the possible sentiment classes including positive, negative, and neutral. Every probability value connotes the probability, or the likelihood, that the message in point falls under the particular sentiment type. As the ultimate sentiment designation of the message, a choice is made for the best-probability class. This probabilistic formulation is not only able to support the correct classification, but also support confidence based decision making, which is useful in real time monitoring and downstream aggregation. All in all, the sentiment classification model can be useful to scale representations at the token level to sentiment probability, by tradeoffs between expressiveness and computation cost to be deployed in real-time.

3.5. Streaming Aggregation

Streaming aggregation is crucial to converting predictions of individual sentiment, to time-sensitive useful insights that capture dynamics in opinion formation. In sentiment monitoring using a real-time sentiment monitoring system, every processed social media message is given a numerical

ISSN: 3071-3498
PP: 01-13.

sentiment score derived through the results of the sentiment classification model. Special attention should be paid to individual scores, which give finer information, but they can be too changing to be interpreted separately. To solve this, the sentiment scores are added together in sliding time windows, which allow the system to get temporal trends and moving short term changes in the data stream. A sliding time window is a theoretical time period within which a message is constantly shifting forward as additional messages come. In every window, the sentiment scores that are made in the interval are summed up. The mean sentiment score of a window is a calculation of all the sentiment scores (with all messages in a window) and then the outcome divided with the number of messages that have been seen in the same window. Demonstrating the calculation in intuitive terms, this would derive one representative sentiment value which would portray the general emotive tone of the stream at that particular point in time. The system gives an actual time representation of sentiment development by constantly changing this value with the movement of the window. Large-scale streaming environments have a number of benefits with this aggregation strategy. First, it eliminates outliers or anomalous messages in noisy data, which creates more stable sentiment signal. Second, sliding windows allow fine-temporal resolution and thus, the system is capable of identifying sudden changes of sentiment triggered by breaking events or trending topics. Also, window parameters (duration, step-size) can be adjusted to provide responsiveness and stability (depending on application demands). All in all, streaming aggregation can transform high-frequency sentiment prediction into insights in the form of interpretable time-based summaries, which is an essential element of real-time analytics, visualization, and decision support systems.

4. RESULTS AND DISCUSSION

4.1. Experimental Setup

The proposed sentiment monitoring framework was experimentally tested through massive data sets of social media that were processed in a distributed streaming architecture. The datasets are formed of the continuously generated textual information, which is obtained on the public social media streams and covers a variety of topics, linguistic styles, and sentiment distributions. In order to replicate real-time conditions, the data were fed continuously as a stream of data as opposed to a batch of data to allow testing the behavior of the system at high-velocity and high-volume workload. The experimental environment was run on a distributed computing cluster that had provision to support parallel processing, fault tolerance, and scalable resource allocation as is typical of a real we deployment situation. The streaming architecture unites the data ingestion, preprocessing, and sentiment classification modules and the aggregation modules into a single pipeline. All the modules can be run simultaneously which implies that the system has a low end to end delay and can be used to process incoming messages. The sentiment classification model had been trained on annotated sentiment data and was part of the streaming pipeline to make real-time inference. Aggregation mechanisms that were characteristic of windows, were used to summarize the sentiment trends over time, thus allowing constant performance monitoring to be done during testing. Three key metrics were used to assess system performance, and they were accuracy, latency and throughput. The accuracy is used to evaluate the accuracy of models on sentiment predictions by comparing the output of the model with ground-truth sentiment labels of benchmark datasets. Latency is the amount of time that it takes between the arrival of the message at the ingestion layer and the production of the corresponding sentiment output and indicates how responsive the system is. The number of messages that are handled during a specific period of time is known as throughput that is used to determine how practical and scalable the streaming framework is. The combination of the metrics analyzed at the same time in the context of the experiment will give a comprehensive assessment of the efficiency in the analysis and real-time performance, which will prove the correctness of using the proposed framework in the case of large-scale sentiment monitoring.

4.2. Comparative Analysis

Compared analysis is an evaluation of the effectiveness of various sentiment analysis methods in accuracy, latency efficiency, and scalability, showing the trade-offs associated with the existing sentiment monitoring systems to target a particular model in real time. Both methods have different strengths and weaknesses according to the means of their computation and the representational ability.

Table 1: Comparative Analysis

Method	Accuracy (%)	Latency Efficiency (%)	Scalability (%)
Lexicon-Based	68.4	90	92
SVM (TF–IDF)	75.9	70	68
LSTM	82.3	55	65
Transformer-Based	86.7	50	60

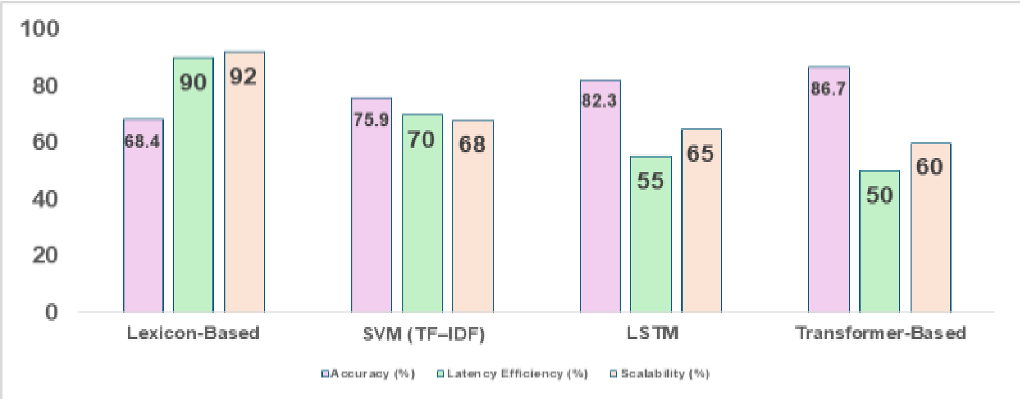

Fig 3 - Comparative Analysis

4.2.1. Lexicon-Based Methods

As a method of sentiment analysis based on lexicon, lexicon-based sentiment analysis is moderately accurate as it is dependent on preset sentiment dictionaries and rule-based aggregation. Although this method has lower predictive performance than learning-based methods, the method has very high latency efficiency and scalability. Its light-weighted computation and lack of model inference render it extremely applicable to real-time systems with explicit resource demands or with the demand to deploy swiftly about extensive quantities of data.

4.2.2. SVM with TF–IDF Features

TFIDF representations in Support Vector Machine models lead to a significant improvement in accuracy since prediction patterns of sentiment are learned using labeled data. Nevertheless, such a benefit comes over the cost of lower latency efficiency and scaling. This additional computational cost of feature vectors construction and kernel-based classification places further computational load on a kernel-based system, and SVM-based systems are therefore less responsive in high-throughput streaming applications than lexicon-based systems.

4.2.3. LSTM-Based Models

Sentiment classifier LSTMs are highly more accurate, in effect capturing text sequential dependencies and context using text. Their propensity to resolve both long-range relationships

ISSN: 3071-3498
PP: 01-13.

allows understanding sentiments more refinement. However, LSTM networks are not as efficient on latency and have moderate scalability due to the periodicity of its implementation which raises inference time and computation scale, especially in processing large amounts of streaming data in real-time.

4.2.4. Transformer-Based Models

Transformer models are the most accurate of all the tested models because they have mechanisms to focus on objects and embed them in the environment. These models are very effective in processes of dealing with complex structures in linguistic expressions and in nuances of sentiment. Nonetheless, they are characterized by a large parameter size and self-attention computation that make them less scalable with respect to latency and scalability than simpler models. Consequently, optimisation approaches are frequently needed to compromise accuracy with real-time performance requirements in transformer-based approaches.

4.3. Discussion

The experimental findings also emphasize a trade-off between the accuracy of sentiment classification and the real-time system performance especially where transformer-based models have been used. Transformer architectures are more accurate since they rely on self-attention and contextual embeddings to dynamically learn complicated semantic relations in text. This would enable them to manage more finer linguistic phenomena like contextual polarity shifts, sarcasm and implicit sentiment expression better than traditional or recurrent models. Nonetheless, such advantages are met by a significant computational overhead in terms of memory utilization and slower inference time, which is undesirable in real time high-throughput setting in streaming. This overhead is a major deployment burden in real-time sentiment monitoring systems where predictive accuracy is not as important as timely insights. The efficiency of the system in responding to a quickly-developing event would be diminished because high latency may result in delayed sentiment updates. In order to deal with this challenge, more focus has been laid on lightweight neural architecture and model optimization strategies. Compression methods like model compression, parameter pruning, quantization, and knowledge distillation can lead to a high shrinkage of models as well as computational costs and effort, with little loss in the original predictive performance. These methods are also endowed with the ability to make inference faster and consume fewer resources by eliminating redundant parameters and simplifying numerical representations, thus making highly sophisticated sentiment models more appropriate to real-time applications. Also, hybrid systems design, a mix between lightweight models to filter initial and transformer based models to analyze the details provides a promising compromise. This type of architecture can devote speed to most received messages and apply more elaborate models to high-impact or uncertain messages. On balance, the discussion highlights that it is essential to determine the optimal methods of achieving accuracy, latency and scalability in designing real-time sentiment analysis systems. As part of future work, it would be prudent to consider co-optimization of model structures and streaming systems so that, rich contextual knowledge may be attained without compromising on the real time performance.

5. CONCLUSION

The present paper has provided an in-depth analysis of real-time sentiment tracking on social media streams, both in terms of the method development and system-level issues that need to be addressed to implement the system in the real world. In a comprehensive overview of the existing

ISSN: 3071-3498
PP: 01-13.

sentiment analysis methods, including lexicon-based and classic machine learning analysis methods, deep learning and transformer-based deep learning, the article showed the trend towards more context-sensitive and accurate sentiment analysis techniques. The given modular streaming framework showed that the data ingestion, preprocessing, feature representation, sentiment classification and aggregation can be efficiently allied to help in scalable and low-latency sentiment analytics in dynamic social media settings. The experimental test indicated the advantages and disadvantages of various sentiment classification models in real-time environments. The findings were also very clear; deep learning structures especially transformer-based models come up with significantly larger accuracy in capturing rich contextual and semantic data. Nevertheless, this enhanced predictive accuracy is introduced by performance in the form of high computational complexity, inference latency, and scalability when handling large streaming loads. Lexicon-based and classical machine learning are less accurate, but are still desirable in an application with strict real-time requirements because they are efficient and simplistic. These results highlight the importance of not only choosing sentiment analysis models depending on their accuracy but also the needs of the application. The research also showed that architectural design is very important in addressing the gap between the performance in terms of analytics and the responsiveness of the system. With distributed streaming pipelines, window-based aggregation tactics and optimized preprocessing mechanisms, sentiment tracking at scale is possible. In addition, the discussion also indicated that lightweight neural models, hybrid architecture and model compression methods have a possibility of reducing challenges related to latency experienced with complex deep learning models, thus rendering them more practical in real-time applications. To continue on to the future, this work shows some promising directions of research. The sentiment analysis that is related to multiple languages is an open confounder since the social media is full of multi-lingual communities with different cultural representations of the sentiment. Moreover, explainable sentiment models should be developed to enhance transparency and trust especially when used in high stakes scenarios, like in public policy analysis and market intelligence. Such ethical aspects as the privacy of data used, the elimination of bias, and responsible use of data in the social media should also be of a priority. Since social media is still affecting the masses and shaping the societal debate, real-time sentiment monitoring can still provide an essential instrument in comprehending the mass behavior and providing prompt and informed decision-making.

REFERENCES

[1] B. Liu, Sentiment Analysis and Opinion Mining, Morgan & Claypool Publishers, 2012.
[2] M. Hu and B. Liu, “Mining and summarizing customer reviews,” in Proc. ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, Seattle, WA, USA, 2004, pp. 168–177.
[3] A. Esuli and F. Sebastiani, “SentiWordNet: A publicly available lexical resource for opinion mining,” in Proc. LREC, Genoa, Italy, 2006, pp. 417–422.
[4] P. D. Turney, “Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews,” in Proc. ACL, Philadelphia, PA, USA, 2002, pp. 417–424.
[5] B. Pang, L. Lee, and S. Vaithyanathan, “Thumbs up? Sentiment classification using machine learning techniques,” in Proc. EMNLP, Philadelphia, PA, USA, 2002, pp. 79–86.
[6] A. McCallum and K. Nigam, “A comparison of event models for Naïve Bayes text classification,” in Proc. AAAI Workshop on Learning for Text Categorization, Madison, WI, USA, 1998, pp. 41–48.
[7] T. Joachims, “Text categorization with support vector machines: Learning with many relevant features,” in Proc. ECML, Chemnitz, Germany, 1998, pp. 137–142.
[8] C. Cortes and V. Vapnik, “Support-vector networks,” Machine Learning, vol. 20, no. 3, pp. 273–297, 1995.
[9] R. Socher et al., “Recursive deep models for semantic compositionality over a sentiment treebank,” in Proc. EMNLP, Seattle, WA, USA, 2013, pp. 1631–1642.

ISSN: 3071-3498
PP: 01-13.

[10] Y. Kim, "Convolutional neural networks for sentence classification," in Proc. EMNLP, Doha, Qatar, 2014, pp. 1746–1751.
[11] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997.
[12] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in Proc. NAACL-HLT, Minneapolis, MN, USA, 2019, pp. 4171–4186.
[13] A. Vaswani et al., "Attention is all you need," in Proc. NIPS, Long Beach, CA, USA, 2017, pp. 5998–6008.
[14] N. Garg et al., "Streaming sentiment analysis using Apache Kafka and Spark," International Journal of Computer Applications, vol. 178, no. 7, pp. 1–6, 2019.