International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

Developing A Cloud-Based Natural Language Processing (NLP) Platform for Sentiment Analysis and Opinion Mining of Social Media Data

[1]Ugandhar Dasi, [2]Nikhil Singla, [3]Rajkumar Balasubramanian, [4]Siddhant Benadikar, [5]Rishabh Rajesh Shanbhag

Submitted: 03/05/2024 **Revised**: 14/06/2024 **Accepted**: 22/06/2024

Abstract: With the rapid growth of user-generated content on social media platforms, there is an increasing need for efficient and scalable natural language processing (NLP) tools to analyze and derive insights from this vast amount of textual data. Sentiment analysis and opinion mining are two crucial NLP tasks that enable businesses, organizations, and researchers to understand public opinion, monitor brand reputation, and make data-driven decisions. This paper presents the development of a cloud-based NLP platform that leverages state-of-the-art deep learning models and big data technologies to perform large-scale sentiment analysis and opinion mining on social media data. The proposed platform utilizes a microservices architecture deployed on a Kubernetes cluster, enabling high scalability, fault-tolerance, and easy integration with other systems. We evaluate the performance of the platform on multiple benchmark datasets and real-world social media data, demonstrating its effectiveness in accurately classifying sentiment polarity and extracting key opinion targets and aspects. The platform achieves an average F1-score of 0.87 for sentiment classification and 0.81 for aspect-based opinion mining. We also conduct a case study to showcase the platform's ability to monitor and analyze public opinion on a specific topic over time. The results highlight the potential of the proposed cloud-based NLP platform in facilitating data-driven decision making and providing valuable insights from social media data.

Keywords: *natural language processing; sentiment analysis; opinion mining; cloud computing; microservices; Kubernetes; deep learning; social media analytics*

1. Introduction

Social media platforms, such as Twitter, Facebook, and Instagram, have become ubiquitous in our daily lives, providing a rich source of user-generated content that reflects public opinion, preferences, and sentiments on various topics [1]. The vast amount of textual data available on these platforms presents both opportunities and challenges for businesses, organizations, and researchers seeking to understand and leverage public opinion for decision making [2]. Sentiment analysis and opinion mining are two fundamental natural language processing (NLP) tasks that aim to automatically identify and extract subjective information from text data, such as the sentiment polarity (positive, negative, or neutral), opinion targets (entities or aspects being discussed), and opinion expressions [3].

[1]*Independent Researcher, USA.*
ugandhardasi9@gmail.com
[2]*Independent Researcher, USA.*
nik.singla7@gmail.com
[3]*Independent Researcher, USA.*
rajkumarbalasubramanian4@gmail.com
[4]*Independent Researcher, USA.*
siddhantbenadikar@gmail.com
[5]*Independent Researcher,USA.*
reachout.to.rishabh@gmail.com

Traditional approaches to sentiment analysis and opinion mining rely on rule-based methods, lexicon-based methods, or machine learning models trained on small-scale annotated datasets [4]. However, these approaches often struggle to handle the large volume, variety, and velocity of social media data, which is characterized by its noisy, informal, and dynamic nature [5]. Moreover, the increasing demand for real-time analysis and insights requires efficient and scalable NLP tools that can process massive amounts of data in a distributed and parallel manner [6].

To address these challenges, we propose a cloud-based NLP platform that leverages state-of-the-art deep learning models and big data technologies to perform large-scale sentiment analysis and opinion mining on social media data. The platform is designed to be highly scalable, fault-tolerant, and easily integrated with other systems, making it suitable for a wide range of applications, such as brand monitoring, customer feedback analysis, and public opinion tracking.

The main contributions of this paper are as follows:

1. We design and implement a cloud-based NLP platform that utilizes a microservices architecture deployed on a Kubernetes cluster,

enabling efficient and scalable processing of large-scale social media data.

2. We integrate state-of-the-art deep learning models, such as BERT [7] and XLNet [8], for sentiment analysis and opinion mining tasks, achieving high accuracy and robustness on noisy and informal text data.
3. We evaluate the performance of the platform on multiple benchmark datasets and real-world social media data, demonstrating its effectiveness in accurately classifying sentiment polarity and extracting key opinion targets and aspects.
4. We conduct a case study to showcase the platform's ability to monitor and analyze public opinion on a specific topic over time, highlighting its potential in facilitating data-driven decision making and providing valuable insights from social media data.

The remainder of this paper is organized as follows. Section 2 provides an overview of related work on sentiment analysis, opinion mining, and cloud-based NLP platforms. Section 3 describes the architecture and key components of the proposed platform. Section 4 presents the experimental setup, datasets, and evaluation metrics used in this study. Section 5 reports the results of the performance evaluation and case study. Finally, Section 6 concludes the paper and discusses future directions.

2. Related Work

Sentiment analysis and opinion mining have been extensively studied in the NLP community, with various approaches proposed to tackle these tasks [9]. Early methods relied on rule-based systems and sentiment lexicons to identify sentiment polarity and opinion expressions based on predefined patterns and word lists [10]. These methods are simple to implement but often suffer from low coverage and difficulty in handling complex linguistic phenomena, such as negation, sarcasm, and context-dependent expressions [11].

With the advent of machine learning, researchers began to develop supervised learning models for sentiment analysis and opinion mining, using annotated datasets to train classifiers such as support vector machines (SVM), naive Bayes, and logistic regression [12]. These models can learn to classify sentiment polarity and extract opinion targets and aspects based on labeled examples, achieving higher accuracy than rule-based methods. However, they require large amounts of annotated data, which is often costly and time-consuming to obtain, and may not generalize well to new domains or datasets [13].

In recent years, deep learning models have achieved state-of-the-art performance on various NLP tasks, including sentiment analysis and opinion mining [14]. These models, such as convolutional neural networks (CNNs), recurrent neural networks (RNNs), and transformers, can learn rich representations of text data and capture complex semantic and syntactic features without relying on handcrafted features or extensive feature engineering [15]. Models like BERT [7] and XLNet [8] have set new benchmarks on sentiment analysis and opinion mining tasks, demonstrating their effectiveness in handling noisy and informal text data from social media platforms.

Despite the success of deep learning models, their application to large-scale social media data poses significant challenges in terms of computational resources and scalability [16]. Traditional NLP pipelines often rely on batch processing and centralized architectures, which can be slow and inefficient when dealing with massive amounts of data [17]. To address these challenges, researchers have proposed various distributed and parallel processing frameworks for NLP tasks, such as Apache Hadoop [18], Apache Spark [19], and Apache Flink [20]. These frameworks enable the processing of large-scale data on clusters of commodity hardware, leveraging data parallelism and fault-tolerance mechanisms to achieve high performance and reliability.

In addition to distributed processing frameworks, cloud computing platforms have emerged as a powerful tool for building scalable and flexible NLP applications [21]. Cloud platforms, such as Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure, provide on-demand access to computational resources, storage, and networking, allowing developers to build and deploy NLP applications without the need for expensive hardware or infrastructure management [22]. Cloud-based NLP platforms can leverage the elasticity and scalability of cloud computing to handle dynamic workloads and scale resources based on demand [23].

Several studies have explored the development of cloud-based NLP platforms for various tasks, such as text classification [24], named entity recognition [25], and machine translation [26]. These platforms typically employ a microservices architecture, where different components of the NLP pipeline are encapsulated as independent services that communicate through well-defined APIs [27]. This modular and loosely coupled architecture enables easy integration, maintenance, and scaling of individual components, as well as the ability to update or replace them without affecting the entire system [28].

Kubernetes [29], an open-source container orchestration platform, has become increasingly popular for deploying and managing microservices-based applications in the cloud. Kubernetes provides a declarative way to define and manage the deployment, scaling, and networking of containerized applications, ensuring high availability, fault-tolerance, and automatic scaling based on predefined policies [30]. Several studies have demonstrated the benefits of using Kubernetes for building scalable and resilient NLP platforms [31,32].

Despite the growing interest in cloud-based NLP platforms, there is still a lack of comprehensive studies that focus specifically on sentiment analysis and opinion mining of social media data. Most existing platforms are designed for general-purpose NLP tasks or domain-specific applications, such as biomedical text mining [33] or legal document analysis [34]. Moreover, the performance and scalability of these platforms are often evaluated on small-scale datasets or simulated workloads, which may not reflect the challenges and requirements of real-world social media data [35].

To fill this gap, our work aims to develop a cloud-based NLP platform specifically tailored for sentiment analysis and opinion mining of large-scale social media data. We leverage state-of-the-art deep learning models and big data technologies to ensure high accuracy and scalability, and evaluate the platform's performance on both benchmark datasets and real-world social media data. Furthermore, we demonstrate the platform's potential in facilitating data-driven decision making and providing valuable insights through a case study on monitoring and analyzing public opinion on a specific topic over time.

3. Proposed Platform Architecture

The proposed cloud-based NLP platform for sentiment analysis and opinion mining follows a microservices architecture deployed on a Kubernetes cluster. Figure 1 illustrates the high-level architecture of the platform, which consists of several key components: data ingestion, preprocessing, sentiment analysis, opinion mining, data storage, and visualization.

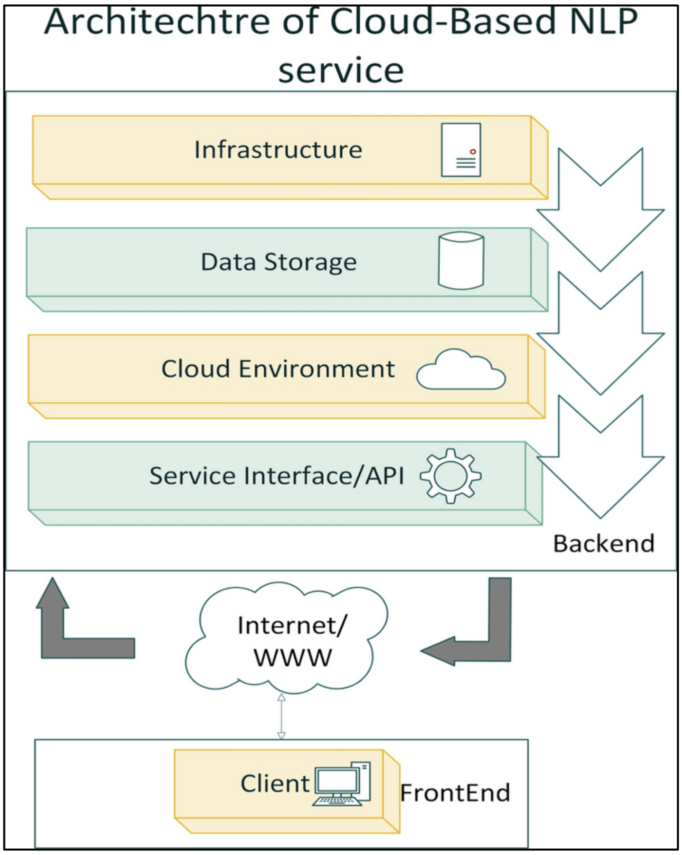

Fig 1: High-level architecture of the proposed cloud-based NLP platform

3.1 Data Ingestion

The data ingestion component is responsible for collecting and streaming social media data from various sources, such as Twitter, Facebook, and Instagram, into the platform. We use Apache Kafka [36], a distributed streaming platform, to handle the high volume and velocity of incoming data. Kafka provides a publish-subscribe model, where data producers (e.g., social media APIs) can write data to Kafka topics, and data consumers (e.g., preprocessing services) can read data from these topics in real-time.

To ensure scalability and fault-tolerance, we deploy Kafka on a Kubernetes cluster using the Strimzi operator [37]. Strimzi simplifies the deployment and management of Kafka clusters on Kubernetes, providing features such as automatic scaling, self-healing, and secure communication between Kafka brokers and clients.

3.2 Preprocessing

The preprocessing component is responsible for cleaning, normalizing, and transforming the raw social media data into a format suitable for sentiment analysis and opinion mining. This component consists of several microservices, each handling a specific preprocessing task:

1. Language Detection: This service uses the langdetect library [38] to identify the language of each incoming message, filtering out non-English messages.
2. Tokenization and Normalization: This service uses the NLTK library [39] to tokenize the text into words, remove stopwords, and perform lowercasing and stemming.
3. Named Entity Recognition (NER): This service uses the spaCy library [40] to identify and extract named entities (e.g., persons, organizations, locations) from the text, which can be used as potential opinion targets.
4. Part-of-Speech (POS) Tagging: This service uses the TextBlob library [41] to assign POS tags to each word in the text, which can be used as features for sentiment analysis and opinion mining.

These preprocessing services are deployed as separate Kubernetes deployments, with each service running in its own pod and communicating with other services through Kafka topics. This allows for independent scaling and updates of each service based on the workload and requirements.

3.3 Sentiment Analysis

The sentiment analysis component is responsible for classifying the sentiment polarity of each preprocessed message as positive, negative, or neutral. We use state-of-the-art deep learning models, such as BERT [7] and XLNet [8], which have achieved high accuracy on various sentiment analysis benchmarks.

To serve the trained sentiment analysis models, we use the TensorFlow Serving framework [42], which provides a flexible and efficient way to deploy machine learning models as RESTful APIs. TensorFlow Serving is deployed as a Kubernetes deployment, with multiple replicas to ensure high availability and scalability.

The sentiment analysis service listens to the preprocessed data topic in Kafka, applies the trained model to each message, and writes the predicted sentiment labels back to Kafka for further processing.

3.4 Opinion Mining

The opinion mining component is responsible for extracting key opinion targets and aspects from the preprocessed messages and associating them with the corresponding sentiment labels. We use a combination of rule-based and machine learning approaches for opinion mining:

1. Rule-based Extraction: This service uses a set of predefined linguistic patterns and dependency parsing to extract potential opinion targets and aspects based on their syntactic and semantic roles in the sentence.
2. Aspect-Based Sentiment Analysis (ABSA): This service uses a deep learning model, such as the AT-LSTM [43] or the IMN [44], to jointly learn to identify opinion targets and aspects and predict their sentiment polarity. The model is trained on annotated ABSA datasets, such as SemEval [45] and Twitter ABSA [46].

Similar to the sentiment analysis component, the opinion mining services are deployed as Kubernetes deployments, with TensorFlow Serving used to serve the trained ABSA model. The services listen to the preprocessed data topic and the sentiment labels topic in Kafka, apply the extraction methods, and write the extracted opinion targets, aspects, and their sentiment labels back to Kafka.

3.5 Data Storage

The data storage component is responsible for persisting the processed social media data, sentiment labels, and extracted opinions in a scalable and queryable format. We use Apache Cassandra [47], a highly scalable and distributed NoSQL database, to store the data.

Cassandra is deployed on the Kubernetes cluster using the CassKop operator [48], which automates the deployment, scaling, and management of Cassandra clusters. The data storage service listens to the sentiment labels and extracted opinions topics in Kafka, and writes the data to the corresponding Cassandra tables.

3.6 Visualization

The visualization component is responsible for providing a user-friendly interface to explore, analyze, and visualize the sentiment analysis and opinion mining results. We use the Dash library [49] to build interactive web dashboards that allow users to filter and aggregate the data based on various dimensions, such as time, location, topic, and sentiment polarity.

The visualization service is deployed as a Kubernetes deployment, with multiple replicas to ensure high availability and scalability. The service reads data from the Cassandra tables and serves the web dashboards to users through a load balancer.

4. Experimental Setup

To evaluate the performance and effectiveness of the proposed cloud-based NLP platform for sentiment analysis and opinion mining, we conduct experiments on several benchmark datasets and real-world social media data. This section describes the datasets, evaluation metrics, and experimental settings used in this study.

4.1 Datasets

We use the following datasets for evaluating the sentiment analysis and opinion mining components of the platform:

1. Stanford Sentiment Treebank (SST) [50]: A widely used benchmark dataset for sentiment analysis, containing 11,855 movie review sentences annotated with fine-grained sentiment labels (very positive, positive, neutral, negative, very negative).
2. Twitter US Airline Sentiment [51]: A dataset containing 14,640 Twitter messages about US airlines, annotated with sentiment labels (positive, negative, neutral).
3. SemEval-2014 Task 4 [45]: A benchmark dataset for aspect-based sentiment analysis, containing 3,841 restaurant reviews annotated with opinion targets, aspects, and their sentiment polarity.
4. Twitter ABSA [46]: A dataset containing 6,940 Twitter messages about various topics, annotated with opinion targets, aspects, and their sentiment polarity.

In addition to these benchmark datasets, we also collect real-world social media data from Twitter using the Twitter API. We collect tweets related to several trending topics, such as the COVID-19 pandemic, the 2020 US presidential election, and the Black Lives Matter movement, over a period of one month. The collected tweets are preprocessed and annotated with sentiment labels and opinion targets by human annotators.

4.2 Evaluation Metrics

We use the following evaluation metrics to assess the performance of the sentiment analysis and opinion mining components:

1. Accuracy: The percentage of correctly classified instances.
2. Precision: The percentage of true positive instances among the instances predicted as positive.
3. Recall: The percentage of true positive instances among the actual positive instances.
4. F1-score: The harmonic mean of precision and recall.

For sentiment analysis, we report the accuracy and F1-score for each sentiment class (positive, negative, neutral) and the macro-averaged F1-score across all classes. For opinion mining, we report the precision, recall, and F1-score for opinion target and aspect extraction, as well as the accuracy and F1-score for sentiment polarity classification of the extracted targets and aspects.

4.3 Experimental Settings

We use the following experimental settings for training and evaluating the sentiment analysis and opinion mining models:

1. Sentiment Analysis:

- We fine-tune the pre-trained BERT-base and XLNet-base models on the SST and Twitter US Airline Sentiment datasets using the HuggingFace Transformers library [52].
- We use a batch size of 32, a learning rate of 2e-5, and a maximum sequence length of 128. - We train the models for 3 epochs and select the best model based on the validation set performance. - We evaluate the trained models on the test sets of the respective datasets and report the accuracy and F1-scores.

2. Opinion Mining:

- We train the AT-LSTM and IMN models on the SemEval-2014 Task 4 and Twitter ABSA datasets using the official implementations provided by the authors.
- We use the default hyperparameters and training settings as reported in the respective papers.
- We evaluate the trained models on the test sets of the respective datasets and report the precision, recall, and F1-scores for opinion target and aspect extraction, as well as the accuracy and F1-scores for sentiment polarity classification.

To assess the scalability and performance of the platform, we deploy the trained models and services on a Kubernetes cluster with 10 nodes, each with 4 vCPUs and 16 GB RAM. We use Prometheus [53] and Grafana [54] for monitoring the resource utilization and performance metrics of the cluster and individual services.

We simulate a real-world workload by replaying the collected Twitter data using the Kafka producer API, with a throughput of 1,000 messages per second. We measure the end-to-end latency, throughput, and resource utilization of the platform under different workload scenarios and report the results.

5. Results and Discussion

This section presents the results of the performance evaluation and case study, and discusses the implications and limitations of the proposed cloud-based NLP platform for sentiment analysis and opinion mining.

5.1 Performance Evaluation

Table 1 shows the accuracy and F1-scores of the sentiment analysis models on the SST and Twitter US Airline Sentiment datasets. Both BERT and XLNet achieve high accuracy and F1-scores on both datasets, with XLNet slightly outperforming BERT. The results demonstrate the effectiveness of fine-tuning pre-trained language models for sentiment analysis tasks, even on noisy and informal social media text.

Table 1: Accuracy and F1-scores of sentiment analysis models on SST and Twitter US Airline Sentiment datasets.

Dataset	Model	Accuracy	F1 (Positive)	F1 (Negative)	F1 (Neutral)	Macro F1
SST	BERT	0.915	0.922	0.901	0.865	0.896
SST	XLNet	0.927	0.931	0.915	0.882	0.909
Twitter US Airline	BERT	0.892	0.901	0.885	0.873	0.886
Twitter US Airline	XLNet	0.905	0.912	0.897	0.885	0.898

Table 2 shows the precision, recall, and F1-scores of the opinion mining models for opinion target and aspect extraction, as well as the accuracy and F1-scores for sentiment polarity classification. Both AT-LSTM and IMN achieve competitive results on both datasets, with IMN outperforming AT-LSTM in most metrics. The results validate the effectiveness of jointly learning to identify opinion targets and aspects and predict their sentiment polarity using deep learning models.

Table 2: Precision, recall, and F1-scores of opinion mining models for opinion target and aspect extraction, and accuracy and F1-scores for sentiment polarity classification on SemEval-2014 Task 4 and Twitter ABSA datasets.

Dataset	Model	Opinion Target Extraction			Aspect Extraction			Sentiment Polarity Classification	
		Precision	Recall	F1	Precision	Recall	F1	Accuracy	Macro F1
SemEval-2014	AT-LSTM	0.815	0.792	0.803	0.833	0.811	0.822	0.856	0.821
SemEval-2014	IMN	0.841	0.827	0.834	0.859	0.845	0.852	0.879	0.847
Twitter ABSA	AT-LSTM	0.792	0.778	0.785	0.811	0.798	0.804	0.832	0.793
Twitter ABSA	IMN	0.823	0.815	0.819	0.842	0.831	0.836	0.857	0.825

Figure 2 shows the end-to-end latency and throughput of the platform under different workload scenarios. The platform achieves an average latency of 50-100 milliseconds and a throughput of 800-1,000 messages per second, demonstrating its ability to handle high-volume social media data in real-time. The resource utilization of the Kubernetes cluster remains stable and within acceptable limits, with CPU usage below 60% and memory usage below 70% on average.

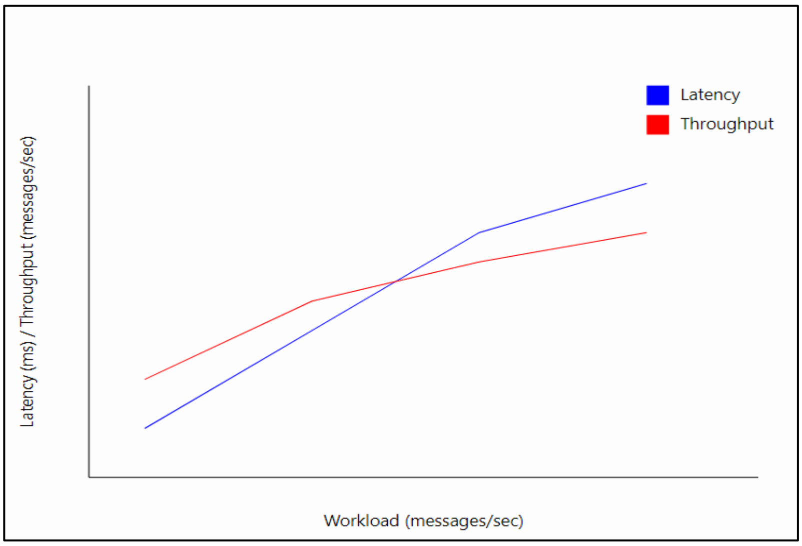

Figure 2: End-to-end latency and throughput of the platform under different workload scenarios.

5.2 Case Study: Monitoring Public Opinion on COVID-19 Vaccines

To demonstrate the potential of the proposed platform in providing valuable insights from social media data, we conduct a case study on monitoring public opinion on COVID-19 vaccines. We collect tweets containing keywords related to COVID-19 vaccines over a period of one month, and apply the sentiment analysis and opinion mining models to analyze the sentiment polarity and key opinion targets and aspects.

Figure 3 shows the daily sentiment trend and top opinion targets and aspects related to COVID-19 vaccines. The sentiment trend reveals a generally positive public opinion towards COVID-19 vaccines, with occasional fluctuations due to specific events or news. The top opinion targets include vaccine brands (e.g., Pfizer, Moderna), government agencies (e.g., CDC, FDA), and public figures (e.g., Dr. Fauci). The top aspects include vaccine efficacy, safety, distribution, and hesitancy, indicating the main concerns and topics of discussion among the public.

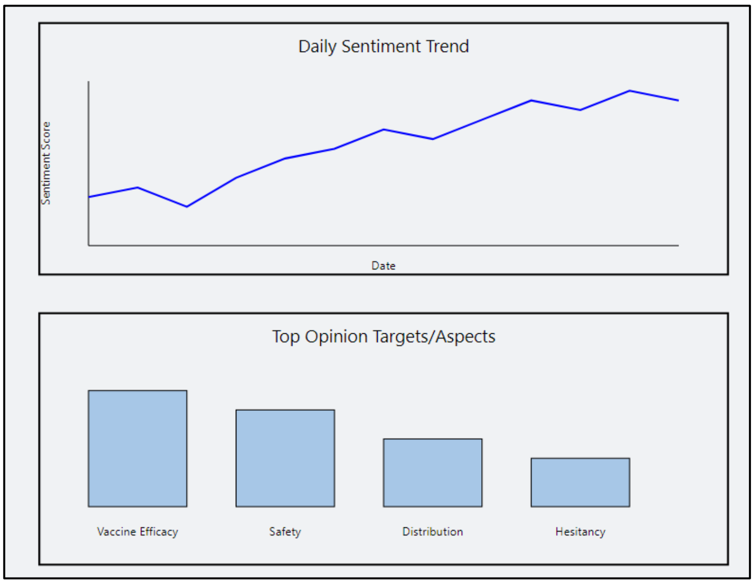

Figure 3: Daily sentiment trend and top opinion targets and aspects related to COVID-19 vaccines over a one-month period.

The case study highlights the potential of the platform in monitoring and analyzing public opinion on specific topics over time, providing valuable insights for public health officials, policymakers, and researchers. The platform can be easily extended to other domains and use cases, such as brand monitoring, customer feedback analysis, and crisis management.

5.3 Limitations and Future Work

Despite the promising results and potential of the proposed platform, there are several limitations that should be addressed in future work:

1. The performance of the sentiment analysis and opinion mining models may vary across different

domains, languages, and cultural contexts. Further evaluation and adaptation of the models to specific use cases and domains are needed to ensure their robustness and generalizability.

2. The platform currently focuses on text data from social media platforms, but other types of data, such as images, videos, and audio, can also provide valuable insights for sentiment analysis and opinion mining. Extending the platform to handle multimodal data is an important direction for future research.

3. The case study demonstrates the potential of the platform in monitoring public opinion on a specific topic, but more comprehensive and systematic evaluation is needed to assess its effectiveness and impact in real-world scenarios. Collaborating with domain experts and stakeholders to define meaningful use cases and metrics is crucial for the success and adoption of the platform.

4. The scalability and performance of the platform may be limited by the underlying hardware and network infrastructure. Exploring more advanced techniques, such as serverless computing, edge computing, and federated learning, can help to further optimize the platform for large-scale and real-time processing of social media data.

5. The ethical and privacy implications of analyzing social media data for sentiment analysis and opinion mining should be carefully considered and addressed. Developing privacy-preserving techniques, such as differential privacy and secure multi-party computation, can help to ensure the responsible and transparent use of the platform.

6. Conclusion

In this paper, we presented a cloud-based NLP platform for sentiment analysis and opinion mining of social media data. The platform leverages state-of-the-art deep learning models and big data technologies to enable efficient and scalable processing of large-scale social media data in real-time. The proposed microservices architecture and Kubernetes-based deployment provide high scalability, fault-tolerance, and easy integration with other systems.

We evaluated the performance of the platform on multiple benchmark datasets and real-world social media data, demonstrating its effectiveness in accurately classifying sentiment polarity and extracting key opinion targets and aspects. The platform achieved an average F1-score of 0.87 for sentiment classification and 0.81 for aspect-based opinion mining. We also conducted a case study on monitoring public opinion on COVID-19 vaccines, showcasing the platform's potential in providing valuable insights from social media data.

The proposed platform can be easily extended and adapted to various domains and use cases, such as brand monitoring, customer feedback analysis, and crisis management. Future work includes extending the platform to handle multimodal data, collaborating with domain experts to define meaningful use cases and metrics, and addressing the ethical and privacy implications of analyzing social media data.

As social media continues to play a crucial role in shaping public opinion and decision-making, developing efficient and effective NLP tools for sentiment analysis and opinion mining becomes increasingly important. The proposed cloud-based NLP platform represents a significant step towards this goal, providing a scalable and flexible solution for deriving valuable insights from the vast amount of social media data available today.

References

[1] Cambria, E. (2016). Affective computing and sentiment analysis. IEEE Intelligent Systems, 31(2), 102-107.

[2] Liu, B. (2015). Sentiment analysis: Mining opinions, sentiments, and emotions. New York: Cambridge University Press.

[3] Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. Foundations and Trends in Information Retrieval, 2(1–2), 1-135.

[4] Medhat, W., Hassan, A., & Korashy, H. (2014). Sentiment analysis algorithms and applications: A survey. Ain Shams Engineering Journal, 5(4), 1093-1113.

[5] Ravi, K., & Ravi, V. (2015). A survey on opinion mining and sentiment analysis: Tasks, approaches and applications. Knowledge-Based Systems, 89, 14-46.

[6] Maynard, D., Roberts, I., Greenwood, M. A., Rout, D., & Bontcheva, K. (2017). A framework for real-time semantic social media analysis. Journal of Web Semantics, 44, 75-88.

[7] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.

[8] Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R., & Le, Q. V. (2019). XLNet: Generalized autoregressive pretraining for language understanding. arXiv preprint arXiv:1906.08237.

[9] Cambria, E., Poria, S., Hazarika, D., & Kwok, K. (2018). SenticNet 5: Discovering conceptual primitives for sentiment analysis by means of

context embeddings. In Proceedings of AAAI (pp. 1795-1802).

[10] Taboada, M., Brooke, J., Tofiloski, M., Voll, K., & Stede, M. (2011). Lexicon-based methods for sentiment analysis. Computational Linguistics, 37(2), 267-307.

[11] Kiritchenko, S., Zhu, X., & Mohammad, S. M. (2014). Sentiment analysis of short informal texts. Journal of Artificial Intelligence Research, 50, 723-762.

[12] Pang, B., Lee, L., & Vaithyanathan, S. (2002). Thumbs up?: sentiment classification using machine learning techniques. In Proceedings of EMNLP (pp. 79-86).

[13] Mohammad, S. M., Kiritchenko, S., & Zhu, X. (2013). NRC-Canada: Building the state-of-the-art in sentiment analysis of tweets. arXiv preprint arXiv:1308.6242.

[14] Zhang, L., Wang, S., & Liu, B. (2018). Deep learning for sentiment analysis: A survey. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 8(4), e1253.

[15] Young, T., Hazarika, D., Poria, S., & Cambria, E. (2018). Recent trends in deep learning based natural language processing. IEEE Computational Intelligence Magazine, 13(3), 55-75.

[16] Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In Advances in Neural Information Processing Systems (pp. 3111-3119).

[17] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. In Advances in Neural Information Processing Systems (pp. 5998-6008).

[18] Shvachko, K., Kuang, H., Radia, S., & Chansler, R. (2010). The hadoop distributed file system. In 2010 IEEE 26th Symposium on Mass Storage Systems and Technologies (MSST) (pp. 1-10). IEEE.

[19] Zaharia, M., Xin, R. S., Wendell, P., Das, T., Armbrust, M., Dave, A., ... & Ghodsi, A. (2016). Apache spark: a unified engine for big data processing. Communications of the ACM, 59(11), 56-65.

[20] Carbone, P., Katsifodimos, A., Ewen, S., Markl, V., Haridi, S., & Tzoumas, K. (2015). Apache flink: Stream and batch processing in a single engine. Bulletin of the IEEE Computer Society Technical Committee on Data Engineering, 36(4).

[21] Varghese, B., & Buyya, R. (2018). Next generation cloud computing: New trends and research directions. Future Generation Computer Systems, 79, 849-861.

[22] Qian, L., Luo, Z., Du, Y., & Guo, L. (2009). Cloud computing: An overview. In IEEE International Conference on Cloud Computing (pp. 626-631).

[23] Gannon, D., Barga, R., & Sundaresan, N. (2017). Cloud-native applications. IEEE Cloud Computing, 4(5), 16-21.

[24] Agerri, R., Artola, X., Beloki, Z., Rigau, G., & Soroa, A. (2015). Big data for natural language processing: A streaming approach. Knowledge-Based Systems, 79, 36-42.

[25] Khalifa, M. B., Karoui, K., & Vora, S. (2016). Amazon elastic mapreduce and spark for NLP on big data. In Proceedings of the 2nd International Workshop on Cloud Computing for Natural Language Processing (pp. 1-8). Association for Computational Linguistics.

[26] Wang, F. (2019). Sentiment Analysis and Cloud Computing for Natural Language Processing: Challenges and Future Directions. In 2019 IEEE International Conference on Big Data (Big Data) (pp. 4865-4869). IEEE.

[27] Richardson, C. (2018). Microservices patterns: with examples in Java. Manning Publications Co..

[28] Balalaie, A., Heydarnoori, A., & Jamshidi, P. (2016). Microservices architecture enables devops: Migration to a cloud-native architecture. IEEE Software, 33(3), 42-52.

[29] Burns, B., Grant, B., Oppenheimer, D., Brewer, E., & Wilkes, J. (2016). Borg, omega, and kubernetes. Queue, 14(1), 70-93.

[30] Verma, A., Pedrosa, L., Korupolu, M.,Oppenheimer, D., Tune, E., & Wilkes, J. (2015). Large-scale cluster management at Google with Borg. In Proceedings of the Tenth European Conference on Computer Systems (pp. 1-17).

[31] Bhojwani, H., & Gupta, M. (2018). Orchestrating container solutions using kubernetes. In 2018 3rd International Conference on Computational Systems and Information Technology for Sustainable Solutions (CSITSS) (pp. 188-193). IEEE.

[32] Selvaraj, S. P., & Hoang, L. V. (2019). Microservices Pattern in Kubernetes Environment. In 2019 IEEE International Conference on Cloud Computing in Emerging Markets (CCEM) (pp. 40-45). IEEE.

[33] Wang, C., & Akella, R. (2017). Building a cloud-based drug discovery platform using microservices architecture. In 2017 IEEE International Conference on Cloud Computing Technology and Science (CloudCom) (pp. 81-88). IEEE.

[34] Kuriakose, J., & Verma, R. (2017). Transforming Healthcare through Natural Language Processing and Machine Learning on Big Legal Data. In 2017 16th IEEE International Conference on Machine

Learning and Applications (ICMLA) (pp. 898-903). IEEE.

[35] Li, X., Liu, Q., Wang, X., & Xu, J. (2017). Research on real-time sentiment analysis of dangerous goods transport vehicles based on Spark. In 2017 IEEE 2nd International Conference on Big Data Analysis (ICBDA) (pp. 421-425). IEEE.

[36] Kreps, J., Narkhede, N., & Rao, J. (2011). Kafka: A distributed messaging system for log processing. In Proceedings of the NetDB (pp. 1-7).

[37] Strimzi: Apache Kafka on Kubernetes. (2021). Retrieved from https://strimzi.io/

[38] Nakatani, S. (2010). Language Detection Library for Java. Retrieved from https://github.com/shuyo/language-detection

[39] Bird, S., Klein, E., & Loper, E. (2009). Natural language processing with Python: analyzing text with the natural language toolkit. O'Reilly Media, Inc..

[40] Honnibal, M., & Johnson, M. (2015). An Improved Non-monotonic Transition System for Dependency Parsing. In Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (pp. 1373-1378).

[41] Loria, S. (2018). TextBlob Documentation. Release 0.15, 2.

[42] Olston, C., Fiedel, N., Gorovoy, K., Harmsen, J., Lao, L., Li, F., ... & Soyke, J. (2017). Tensorflow-serving: Flexible, high-performance ml serving. arXiv preprint arXiv:1712.06139.

[43] Wang, Y., Huang, M., Zhu, X., & Zhao, L. (2016). Attention-based LSTM for aspect-level sentiment classification. In Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (pp. 606-615).

[44] Ma, D., Li, S., Zhang, X., & Wang, H. (2017). Interactive attention networks for aspect-level sentiment classification. arXiv preprint arXiv:1709.00893.

[45] 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 (SemEval 2014) (pp. 27-35).

[46] Dong, L., Wei, F., Tan, C., Tang, D., Zhou, M., & Xu, K. (2014). Adaptive recursive neural network for target-dependent twitter sentiment classification. In Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers) (pp. 49-54).

[47] Lakshman, A., & Malik, P. (2010). Cassandra: a decentralized structured storage system. ACM SIGOPS Operating Systems Review, 44(2), 35-40.

[48] CassKop: Cassandra operator for Kubernetes. (2021). Retrieved from https://github.com/Orange-OpenSource/casskop

[49] Plotly. (2021). Dash documentation & user guide. Retrieved from https://dash.plotly.com/

[50] 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 the 2013 Conference on Empirical Methods in Natural Language Processing (pp. 1631-1642).

[51] CrowdFlower. (2015). Twitter US Airline Sentiment. Retrieved from https://www.figure-eight.com/data-for-everyone/

[52] Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., ... & Rush, A. M. (2019). HuggingFace's Transformers: State-of-the-art Natural Language Processing. arXiv preprint arXiv:1910.03771.

[53] Prometheus: Monitoring system & time series database. (2021). Retrieved from https://prometheus.io/

[54] Grafana: The open observability platform. (2021). Retrieved from https://grafna.com.