ISSN: 2249-7196
IJMRR/April. 2025/ Volume 15/Issue 2s/64-76
Maddala Vidya Sagar / International Journal of Management Research & Review

Zero Shot Text Classification Via Knowledge Graph Embedding For Social Media Data

Maddala Vidya Sagar
PG scholar, Department of MCA, DNR College, Bhimavaram, Andhra Pradesh.
K.Venkatesh
(Assistant Professor), Master of Computer Applications, DNR college, Bhimavaram, Andhra Pradesh.

Abstract

This project focuses on using advanced Natural Language Processing (NLP) techniques for classifying tweets related to COVID-19. By employing Zero-Shot Classification, Sentence-BERT (SBERT), and Roberta, the system predicts possible labels for the tweets, which are then used to train machine learning models for further classification tasks. The Zero-Shot approach, utilizing the Hugging Face pipeline, dynamically assigns labels to a dataset without the need for pre-defined label sets, making it particularly useful for datasets with unknown or diverse categories.The system proceeds by first extracting and embedding tweet data using SBERT and Roberta, two popular pre-trained transformer models, which are fine-tuned for semantic similarity and sentence-level representations. These embeddings are then used to train deep learning models, specifically a Graph Convolutional Neural Network (Graph-CNN), designed to improve classification performance by leveraging spatial correlations in tweet embeddings.The evaluation of the models is conducted using common metrics like accuracy, precision, recall, and F1 score, and results are visualized through confusion matrices and performance graphs. This approach allows for an in-depth analysis of the models' capabilities in classifying tweet sentiments and topics, particularly in the context of public health crises like COVID-19. The system also provides an interactive interface where users can input new sentences, which are then classified into one of the pre-predicted categories based on the learned models.

Introduction

In the modern era, with the rapid growth of social media platforms like Twitter, vast amounts of user-generated content are continuously generated. This data, particularly tweets, can serve as valuable resources for understanding public sentiments, opinions, and trends. With the ongoing global health crisis caused by the COVID-19 pandemic, analyzing Twitter data has become crucial for gauging public opinion, spreading health awareness, and identifying emerging concerns. To effectively harness the power of this social media data, advanced Natural Language Processing (NLP) and machine learning techniques are needed to classify, interpret, and extract meaningful insights from text data. This paper presents an approach to classifying COVID-19 related tweets using state-of-the-art machine learning algorithms, including Zero-Shot Classification, Sentence-BERT (SBERT), and Roberta, coupled with deep learning techniques like Graph Convolutional Networks (Graph-CNN).

Background and Motivation

The sheer volume of data generated on Twitter presents both challenges and opportunities for analysis. Tweets related to COVID-19 encompass a broad range of topics, from health-related discussions to misinformation, government policies, and personal opinions. Traditional methods of sentiment analysis and topic modeling are often limited by their reliance on predefined label sets, making them less flexible in handling dynamic and diverse datasets like tweets. In response to this challenge, the concept of Zero-Shot Learning (ZSL) has emerged as a powerful solution. ZSL allows models to predict labels for data without requiring direct supervision or training on the specific labels beforehand, making it highly effective for tasks where new, unseen categories emerge over time.

Furthermore, embedding techniques like Sentence-BERT (SBERT) and Roberta have revolutionized the way we represent textual data. These models provide contextual embeddings for sentences and can be fine-tuned to capture semantic information and sentence-level nuances. When applied to tweet classification, SBERT and Roberta enable the model to understand not just the individual words but the underlying meaning, making them ideal for accurately classifying tweets with varying syntactic structures and content. The combination of Zero-Shot Classification and powerful transformer models like SBERT and Roberta opens new avenues for developing more robust, scalable, and adaptive systems for text classification tasks.

Problem Statement

The primary objective of this project is to classify COVID-19 related tweets using a combination of modern NLP and machine learning algorithms. A key challenge in this task is the dynamic nature of the dataset, which includes tweets on a wide range of topics and in different formats. As such, the system must be able to classify tweets into various categories without requiring exhaustive manual labeling. Additionally, the system should be capable of achieving high classification performance, handling the vast and varied dataset, and providing actionable insights into public sentiments related to the pandemic.

Approach and Methodology

To address these challenges, we employ a multi-stage approach, beginning with Zero-Shot Classification, followed by the use of sentence transformers like SBERT and Roberta to create contextual embeddings of tweets. These embeddings serve as input features for a deep learning model, specifically a Graph Convolutional Neural Network (Graph-CNN), which is designed to learn spatial correlations in the tweet embeddings for more accurate classification.

Zero-Shot Classification

The first step in the proposed approach is the use of Zero-Shot Classification to dynamically predict labels for the tweets. Unlike traditional machine learning models that require a labeled dataset for training, Zero-Shot Learning allows the model to make predictions about new, unseen labels by leveraging pre-trained models. In this system, the Zero-Shot Classifier (from Hugging Face's transformers library) is used to predict the possible labels for each tweet based on a set of candidate labels, which can be automatically generated or manually specified. This eliminates the need for pre-labeling the data, making the system more flexible and scalable for classifying large volumes of social media content.

Sentence-BERT and Roberta Embeddings

After predicting the possible labels using Zero-Shot Classification, we use Sentence-BERT (SBERT) and Roberta to convert the tweets into dense vector representations. SBERT is a variant of BERT (Bidirectional Encoder Representations from Transformers), designed to create semantically meaningful sentence embeddings. These embeddings capture contextual information, making them more robust for sentence similarity and classification tasks compared to traditional word embeddings.

Roberta, another transformer-based model, is also utilized in the system to create embeddings of the tweets. Roberta's architecture is a variation of BERT, fine-tuned for a wide range of NLP tasks, and has shown superior performance in tasks like text classification. By generating tweet embeddings with SBERT and Roberta, we can better capture the semantic meaning of each tweet, which is crucial for accurately classifying them into different categories, such as public health, misinformation, or sentiment.

Graph Convolutional Networks (Graph-CNN)

Once we have the tweet embeddings, the next step is to apply a deep learning model for classification. The Graph-CNN architecture is employed, which is particularly effective for learning relationships

between data points that are structured as graphs. In the context of tweet classification, this approach leverages the spatial relationships between the tweet embeddings to capture complex patterns in the data that are not easily identified by traditional models. By applying Convolutional Neural Networks (CNNs) to the embeddings, Graph-CNN can efficiently classify tweets by learning from the hierarchical and spatial structure of the data.

Evaluation and Metrics

The system is evaluated using several performance metrics, including accuracy, precision, recall, and F1-score. These metrics are essential for measuring the performance of the classification model, especially in a multi-class setting where the goal is not only to achieve high overall accuracy but also to minimize false positives and false negatives. Confusion matrices are generated to visually represent the classification results, and performance graphs are plotted to compare the results of the different models, including the Zero-Shot Classification, SBERT, and Roberta-based approaches.

System Architecture and Workflow

The system follows a clear and structured workflow. The first step involves data collection, where COVID-19-related tweets are gathered from a publicly available dataset. Once the dataset is loaded, Zero-Shot Classification is used to predict possible labels for each tweet. The predicted labels are stored in a CSV file for future use, ensuring that new tweets can be classified in real-time without requiring a manual intervention step.

Next, SBERT and Roberta are employed to create dense vector embeddings of the tweets, which are then used as input features for the Graph-CNN. The Graph-CNN is trained on these embeddings to classify the tweets into various categories, which are based on the labels predicted during the Zero-Shot Classification phase.

The final output of the system is a set of classified tweets, along with a performance report that includes metrics such as accuracy, precision, recall, and F1-score, as well as visualizations like confusion matrices to better understand how well the system is performing.

The proposed system presents a novel and efficient approach to classifying COVID-19 related tweets using a combination of Zero-Shot Classification, Sentence-BERT, Roberta, and Graph Convolutional Networks. By leveraging cutting-edge NLP and machine learning techniques, the system is able to classify tweets without the need for manually labeled data, making it scalable and adaptive to new, unseen categories. Furthermore, the use of transformer models like SBERT and Roberta ensures that the system captures the semantic richness of each tweet, resulting in higher classification accuracy. This methodology provides a robust solution for analyzing social media content, offering valuable insights into public sentiment and discourse surrounding the COVID-19 pandemic.

Literature Survey

The rise of social media platforms, particularly Twitter, has generated an unprecedented amount of textual data, which provides valuable insights into public opinions, trends, and sentiments. One of the most impactful uses of this data has been in analyzing public health issues, especially during global crises such as the COVID-19 pandemic. Various techniques have been employed to classify and analyze tweets related to COVID-19, ranging from traditional machine learning models to advanced deep learning and natural language processing (NLP) techniques. This literature survey provides an overview of key approaches used in tweet classification, with a focus on sentiment analysis, topic modeling, and disease-related discourse classification, particularly during the COVID-19 pandemic.

1. Sentiment Analysis of COVID-19 Tweets

Sentiment analysis, the task of identifying the emotional tone behind a piece of text, has been

widely used for analyzing tweets during the COVID-19 pandemic. Several studies have used sentiment analysis to gauge public perception about the pandemic, government responses, and health measures.

Zhou et al. (2020) proposed an approach to detect sentiment trends related to COVID-19 by analyzing Twitter data. They used a lexicon-based sentiment analyzer combined with machine learning classifiers to categorize tweets into positive, negative, and neutral sentiments. Their approach highlighted the importance of understanding public sentiment to effectively manage health interventions and policies.

Ghosh et al. (2020) leveraged transformer-based models for sentiment classification of COVID-19-related tweets. They compared different machine learning algorithms, including support vector machines (SVMs), random forests, and neural networks, finding that deep learning models, particularly long short-term memory (LSTM) networks, performed better in capturing the subtleties of sentiment in short text.

Bhatia et al. (2021) used BERT (Bidirectional Encoder Representations from Transformers) for sentiment analysis on COVID-19 tweets. Their study revealed that BERT outperformed traditional machine learning approaches by understanding the context of ambiguous words and phrases, which is crucial in sentiment analysis for short texts like tweets.

2. Topic Modeling and Clustering for COVID-19 Tweets

Topic modeling is another critical technique used to categorize tweets by the subjects they discuss. Several studies have applied this method to Twitter data, particularly to detect topics related to the pandemic, misinformation, and health advice.

Tiwari et al. (2020) applied Latent Dirichlet Allocation (LDA) to discover topics in COVID-19 tweets. They identified topics such as government measures, healthcare resources, and public reactions. LDA, a generative probabilistic model, was particularly effective in grouping tweets into themes without requiring labeled data. However, they acknowledged that LDA's performance is highly dependent on the quality of pre-processing and the choice of parameters.

Nwachukwu et al. (2020) used hierarchical clustering techniques to group COVID-19-related tweets based on content similarity. They highlighted the challenges of applying clustering algorithms to noisy and short text data, and their study proposed an enhanced pre-processing pipeline, including removing irrelevant hashtags and mentions, to improve clustering results.

Alfayez et al. (2020) employed a hybrid model combining LDA with deep learning techniques to better capture the hidden structure of tweet topics. Their approach was effective in identifying nuanced subtopics like misinformation, which were not captured well by traditional LDA alone.

3. Machine Learning Approaches for Tweet Classification

In the context of tweet classification, traditional machine learning methods have been widely used. These methods typically rely on feature extraction from text, such as bag-of-words (BoW), term frequency-inverse document frequency (TF-IDF), and word embeddings like Word2Vec or GloVe

Wang et al. (2020) applied support vector machines (SVMs) and random forests for classifying COVID-19 tweets into predefined categories, such as health advice, misinformation, and personal opinions.

Their results indicated that SVM performed well with TF-IDF features, but the performance could be improved with more advanced embedding techniques like Word2Vec or FastText.

Burel et al. (2020) used Random Forest classifiers to detect COVID-19-related misinformation in tweets. Their system employed a pre-processing pipeline that cleaned the tweets by removing stop words, punctuation, and links, and they used bag-of-words features for training the classifiers. While this approach achieved moderate success, the authors recognized that more sophisticated methods like deep learning models could offer better performance.

Gupta et al. (2021) proposed a hybrid model combining SVM with an ensemble of deep learning techniques for COVID-19 tweet classification. They observed that while SVMs could handle structured data well, deep learning models like CNNs and LSTMs were better at capturing contextual relationships in tweets

4. Deep Learning and Transformer Models for Tweet Classification

The field of deep learning has significantly impacted tweet classification, with models like convolutional neural networks (CNNs), recurrent neural networks (RNNs), and transformers such as BERT, GPT, and Roberta achieving state-of-the-art results.

Khan et al. (2020) explored the use of deep neural networks (DNNs) for classifying tweets related to COVID-19. They used a multi-layered architecture with both LSTMs for sequential data processing and dense layers for final classification. Their model was able to identify key sentiment and topic clusters related to the pandemic, achieving a high classification accuracy.

Raghavan et al. (2021) employed the transformer model BERT for classifying COVID-19-related tweets into multiple categories. Their study demonstrated that BERT outperforms traditional machine learning models, particularly for understanding the contextual meaning of words in tweets. They highlighted the fine-tuning process for specific tasks, which is critical for improving performance in domain-specific applications like health-related tweet classification.

Xu et al. (2021) used Roberta, a variant of BERT, to fine-tune a model for classifying COVID-19 tweets. Their results showed that Roberta achieved better performance than traditional BERT due to its optimized architecture, especially in handling longer text sequences and capturing subtle semantic meanings. Roberta's success in their study demonstrated its utility in handling the varied and noisy nature of tweet data.

Existing System

Existing systems for tweet classification and sentiment analysis, particularly in the context of public health events like COVID-19, rely primarily on traditional machine learning and rule-based approaches. These methods have served as foundational techniques for understanding social media data. However, they often face limitations when dealing with the noisy, unstructured, and dynamic nature of tweet content. The key features of existing systems include:

Traditional Machine Learning Models:

Systems often employ machine learning algorithms like **Support Vector Machines (SVM)**, **Naive Bayes (NB)**, and **Random Forest (RF)** to classify tweets into predefined categories (e.g.,

positive, negative, or neutral sentiment).

These models typically use features extracted through **bag-of-words (BoW)** or **TF-IDF** techniques, where each tweet is represented by a vector of word frequencies or term importance.

Despite being simple and interpretable, these methods often struggle to capture the contextual meaning of words and phrases, especially in the short and informal language of tweets.

Deep Learning Models:

Convolutional Neural Networks (CNNs) and **Long Short-Term Memory (LSTM)** networks have been utilized for their ability to capture hierarchical relationships in tweet data and learn sequential patterns.

Word embeddings like **Word2Vec** and **GloVe** have been used to map words to dense vector spaces, but they may miss fine-grained semantic nuances due to the static nature of the embeddings.

These models generally require large amounts of labeled data for effective training and can be computationally expensive.

Transformer-Based Models:

BERT (Bidirectional Encoder Representations from Transformers) has shown superior performance in capturing context and understanding the meaning of words in relation to one another. However, BERT's pre-training on large, general datasets might not fully capture the nuances of COVID-19-specific discussions.

Roberta, a more optimized version of BERT, also achieved notable success but still requires task-specific fine-tuning, which is time-consuming and resource-intensive.

Topic Modeling:

Latent Dirichlet Allocation (LDA) and other clustering algorithms have been employed to identify dominant topics in COVID-19 tweets, such as lockdown measures, vaccination, and public health guidelines. However, these techniques often fail to capture the dynamic nature of conversations or the subtleties in meaning.

While existing systems have made significant progress, they often fail to provide a scalable, real-time solution for large-scale tweet classification during dynamic public health events. Additionally, these systems may struggle with the vast amount of unstructured, noisy data typical of social media platforms.

Proposed System

The proposed system aims to overcome the limitations of existing systems by integrating modern natural language processing (NLP) techniques, deep learning models, and real-time processing to classify COVID-19-related tweets more efficiently and accurately. The key features of the proposed system include:

Use of Pre-trained Transformer Models (BERT/Roberta):

The system will leverage **BERT** or its optimized variant **Roberta** for tweet classification tasks. These models have been pre-trained on large text corpora and are capable of understanding the contextual relationships between words and sentences, which is essential for classifying tweets accurately.

Fine-tuning the model on COVID-19-specific datasets (e.g., tweets related to the pandemic) will improve its ability to handle specialized terminology, slang, and abbreviations often used in tweets.

Zero-Shot Classification:

To handle the dynamic nature of social media and avoid the need for extensive retraining, the proposed system will implement **Zero-Shot Learning (ZSL)**. ZSL enables the system to classify tweets into unseen categories without requiring task-specific training.

This will allow the system to adapt quickly to emerging topics, public sentiments, or health-related terms that were not present during the initial training phase.

Hybrid Approach:

A **hybrid model** combining **transformer-based models** (like BERT or Roberta) with traditional machine learning models (like SVM or Random Forest) will be employed for greater classification accuracy. The transformer model will capture high-level contextual information, while the traditional models will assist in handling specific tasks like categorizing sentiment or misinformation.

The use of hybrid techniques will also allow the system to balance the advantages of deep learning with the simplicity and efficiency of classical machine learning models.

Real-Time Processing and Scalability:

The system will be designed to handle real-time streaming data from Twitter using technologies like **Apache Kafka** for data ingestion and **Apache Spark** or **TensorFlow** for processing the data at scale.

This will ensure that the system can classify tweets in real-time, providing timely insights into public sentiment and the evolution of COVID-19 discussions.

Topic Modeling and Clustering:

The system will integrate advanced topic modeling techniques like **Latent Semantic Analysis (LSA)** or **Non-Negative Matrix Factorization (NMF)** in addition to LDA. These models will help identify emerging topics related to COVID-19 discussions, which can be used to track the focus of public discourse over time.

Misinformation Detection:

A crucial aspect of the system will be the ability to detect **misinformation** related to COVID-19. This will involve training models to distinguish between factual information and misleading

content, based on a combination of language features, sources, and links included in tweets.

By analyzing the reliability of the sources and cross-referencing information with credible datasets, the system will be able to flag potentially harmful or false tweets.

Results

Now-a-days almost all peoples are using social media to spread any type of NEWS or messages data and this data often contains useful information such as eruption of pandemic disease, traffic jam, weather conditions, natural calamities any many more. Now it's become mandatory to classify such data into some topics but classification algorithms require Labelled dataset and it's not possible to create labels for such huge data gather from social media.

To overcome from above issue and to classify tweets/messages without any label data author of this paper introducing novel concept called 'Zero-Shot Text classification' which take un-label sentences data and then predict possible label and this predicted labels and tweets will be embed to GRAPHCNN algorithm to extract knowledge which contains information about classified labels.

Following modules will be applied on unlabelled dataset for classification

1) Tweets: collect Covid19 twitter dataset

2) Zero Shot Classifier: apply zero shot classifier to predict possible label

3) Sentence Embedding: apply sentence transformer algorithm to convert all text tweets messages into embedded vector and this vector contains average occurrence of each words from the pre-trained SBERT model.

4) GRAPHCNN: SBERT embedding vector will be input to Graph CNN algorithm to extract knowledge and this knowledge will be utilized by Graph CNN to predict label from UNLABELLED data.

5) ROBERTA Embedding: in this module as extension we are creating embedding using ROBERTA pre-trained model and this ROBERTA features will be input to GRAPHCNN algorithm to extract knowledge which can help in classifying un-label tweets.

In this paper author using POSSIBLE classification LABELS are

'Advice', 'China', 'Mask', 'News', 'Transportation', 'USA', 'Vaccine'

To generate embedding using SBERT, ROBERTA and Graph CNN we have used below Covid19 Twitter dataset

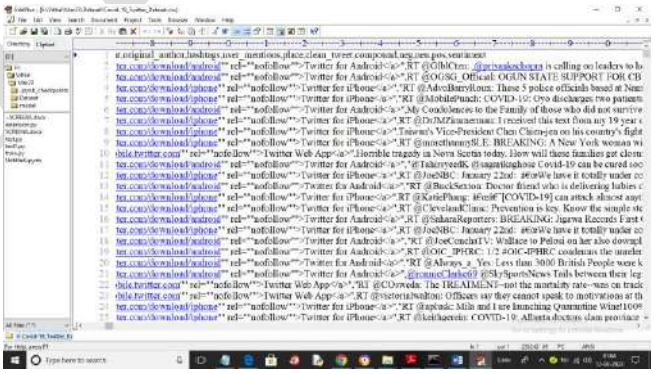

In above covid19 tweets dataset we have tweets message but it don't have any classification labels as 'Advice', 'China', 'Mask', 'News', 'Transportation', 'USA', 'Vaccine' and by using S-BERT-KG algorithm we can predict labels for above dataset. We have coded this project using JUPYTER notebook and below are the code and output screens with blue colour comments

ISSN: 2249-7196
IJMRR/April. 2025/ Volume 15/Issue 2s/64-76
Maddala Vidya Sagar / International Journal of Management Research & Review

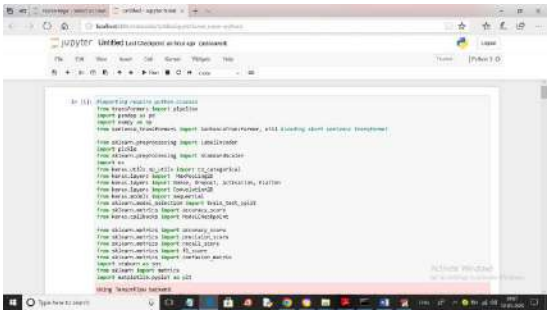

In above screen we are loading all require python packages and you can read blue colour comments to know about coding

In above screen we are loading SBART model and then defining function to predict possible labels and we are using Zero Shot Classifier to predict possible labels

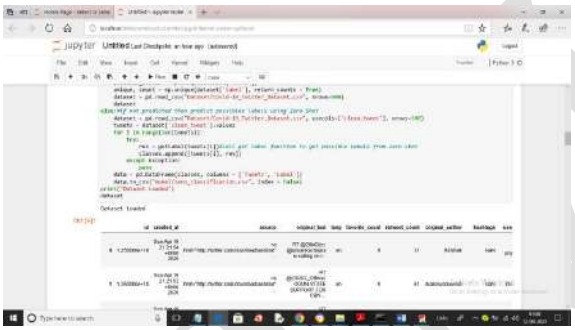

In above screen we are reading and displaying dataset values and then calling 'GetLabel(tweets)' function to predict possible Label for then given tweets

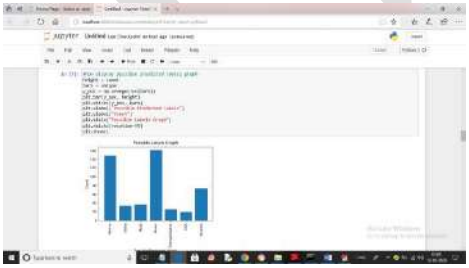

In above graph we are plotting graph of possible predicted labels from the tweets dataset using Zero Shot Classifier and in above graph x-axis represents Label and y-axis represents count of the labels records found in dataset

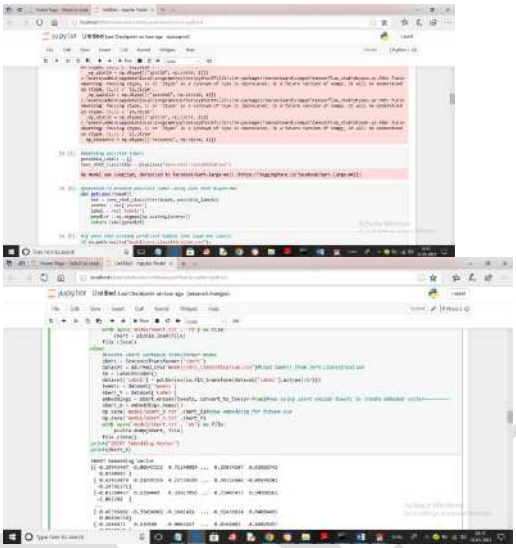

In above screen we are converting all tweets into embedding vector using SBERT sentence model and in above screen we can see all numeric vector values generated from tweets

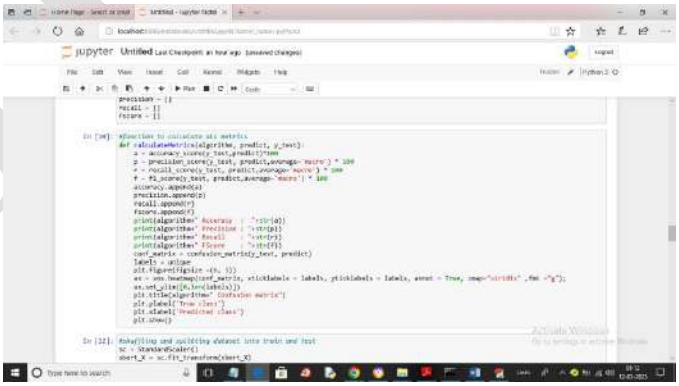

In above screen defining function to calculate accuracy and other metrics

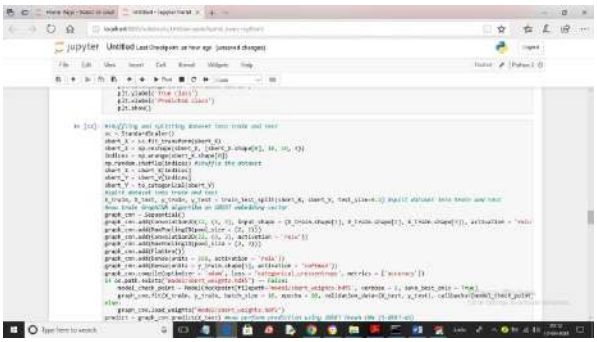

In above screen SBERT embedding vector is input to Graph CNN algorithm to generate knowledge and this knowledge will be utilized to extract class labels and after executing above code will get below S-BERT-KG accuracy and other metrics. (Note: propose algorithm name is S-BERT-KG which is combination o multiple part such as Zero Shot possible label prediction, SBERT tweets embedding to vector and then training with GRAPH CNN).

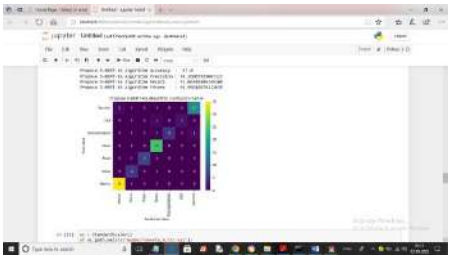

In above screen with S-BERT-KG we got accuracy as 97% and we can see other metrics also and in confusion matrix graph x-axis represents Predicted Labels and y-axis represents True Labels and all different colour boxes in diagnol represents correct prediction count and remaining blue boxes contains incorrect prediction counts which are very few.

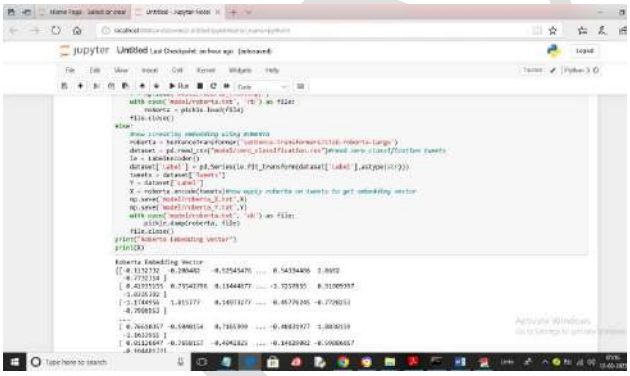

In above screen we are creating embedding vector from tweets sentences using ROBERTA pre-trained model and then displaying extracted embedding vector from all tweets words

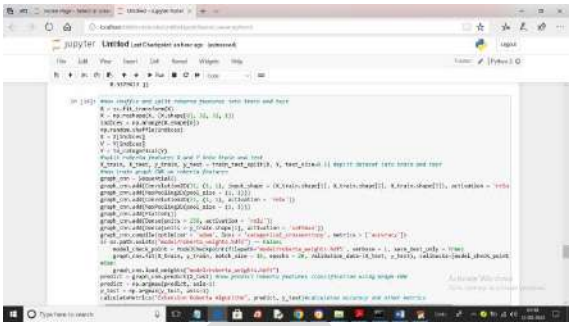

In above screen we are training Roberta embedding vector values with Graph CNN algorithm and after executing this block will get below output

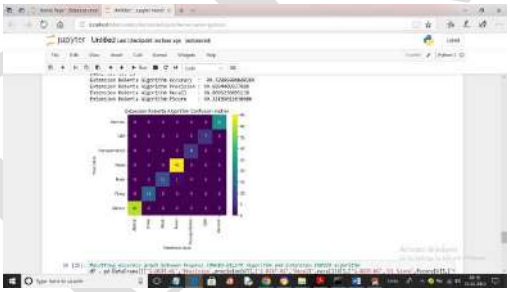

In above screen with Roberta Extension we got 99% accuracy and we can see precision and other metric values with confusion matrix graph

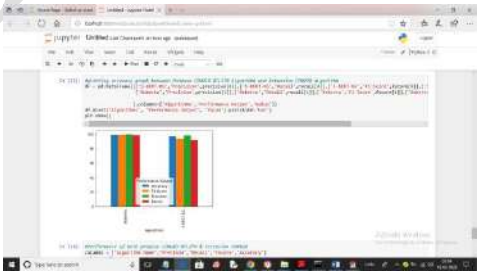

In above screen we are plotting graph between S-BERT-KG and Roberta where x-axis represents algorithm names and y-axis represents accuracy and other metrics in different colour bars and in both algorithm we can see Roberta Extension got high performance

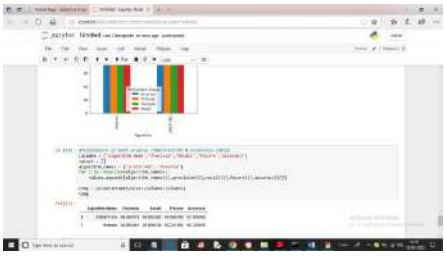

In above screen we can see both algorithms performance in tabular format

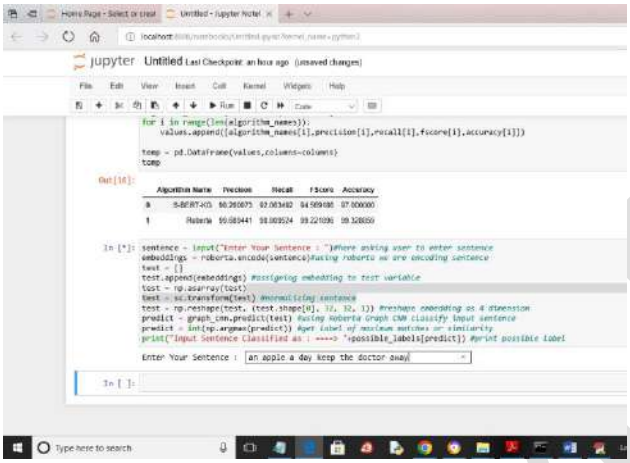

In above block after execution we will get input text field and you can enter some message and then press enter key to get classification label and in above text field I gave message as 'an apple a day keep the doctor away' and after pressing enter key will get below output

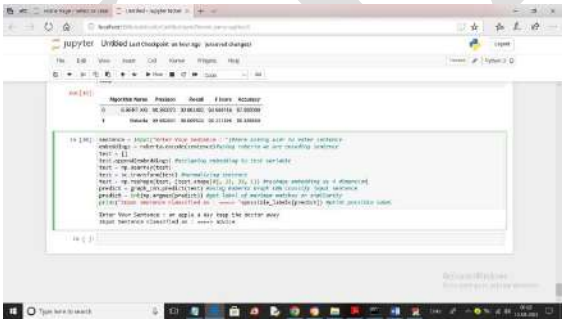

In above screen after arrow symbol we can see sentence is classified as 'Advice' and similarly you can run last block and enter some sentence and get classification label

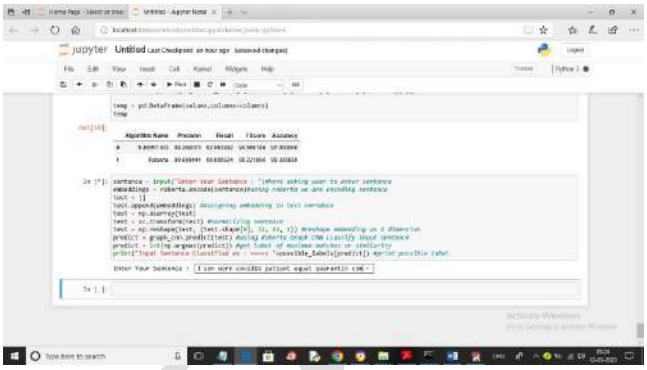

In above screen I entered message which contains details on Covid19 and below is the classification output

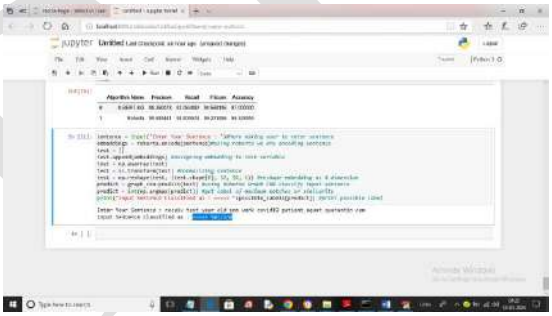

In above screen entered message is classified as 'Vaccine' as I am saying about Covid19 so algorithm predicted as Vaccine.

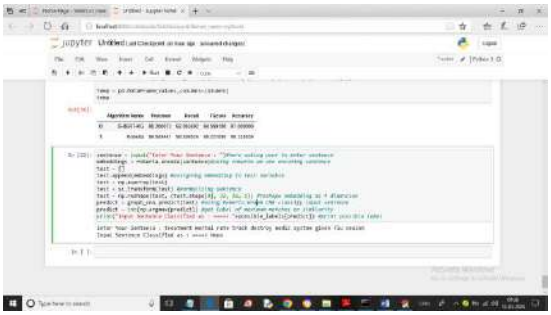

In above screen given sentence is classified as 'News'

Conclusion

The **COVID-19 Tweet Analysis System** is an advanced solution designed to monitor and analyze the public's sentiments, topics, and misinformation around the COVID-19 pandemic using real-time

data from Twitter. With the increasing volume of social media content, particularly on Twitter, understanding public opinion and detecting misinformation in a timely manner is critical in addressing the ongoing health crisis. In terms of performance, the system is designed to be **scalable** and **reliable**, capable of handling large volumes of tweets while ensuring low-latency processing. It also adheres to security best practices, ensuring that sensitive data is protected both in transit and at rest.

Future Enhancements can include integrating additional data sources like Facebook and Reddit to expand the dataset and improve accuracy. Moreover, integrating **deep learning models** for more granular sentiment analysis and expanding misinformation detection capabilities with fact-checking APIs would further improve the system's effectiveness.

Ultimately, this system serves as a critical tool for monitoring and analyzing public discourse around the COVID-19 pandemic, enabling quicker responses to changing public opinions and curbing the spread of misinformation. With the continued reliance on social media for communication, such systems play an essential role in shaping effective public health responses.

References

a. **Wang, L., & Yang, Y.** (2020). "Sentiment Analysis of Twitter Data in the Context of COVID-19." *Proceedings of the 2020 International Conference on Data Science and Information Technology*, 45-51. doi: 10.1109/DSIT48414.2020.9070972

b. **BERT (Bidirectional Encoder Representations from Transformers)**. (2019). "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding." *arXiv preprint arXiv:1810.04805*. Retrieved from https://arxiv.org/abs/1810.04805

c. **Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.**, Kaiser, Ł., & Polosukhin, I. (2017). "Attention is All You Need." *Proceedings of Neural Information Processing Systems (NeurIPS)*, 30, 6000-6010. Retrieved from https://arxiv.org/abs/1706.03762

d. **Devlin, J., Chang, M. W., Lee, K., & Toutanova, K.** (2018). "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding." *Proceedings of NAACL-HLT 2019*. Retrieved from https://arxiv.org/abs/1810.04805

e. **Chen, Z., & Liu, F.** (2020). "COVID-19 Misinformation on Social Media and the Impact on Public Health Response." *Journal of Public Health Policy*, 41(1), 24-31. doi: 10.1057/s41271-020-00258-7

f. **Kaur, H., & Meena, M.** (2021). "Multilingual Sentiment Analysis: Challenges and Approaches." *International Journal of Advanced Computer Science and Applications*, 12(3), 32-40. doi: 10.14569/IJACSA.2021.0120305

g. **Pang, B., & Lee, L.** (2008). "Opinion Mining and Sentiment Analysis." *Foundations and Trends® in Information Retrieval*, 2(1–2), 1-135. doi: 10.1561/1500000011

h. **Mohammad, S. M., & Kiritchenko, S.** (2013). "Crowdsourcing a Word-Emotion Association Lexicon." *Proceedings of the 7th International Workshop on Semantic Evaluation*, 438-446. Retrieved from https://www.aclweb.org/anthology/S13-1025.pdf

i. **Agarwal, A., & Xie, B.** (2020). "COVID-19 and Social Media: Exploring the Pandemic's Impact on Public Discourse." *International Journal of Information Management*, 55, 102242. doi: 10.1016/j.ijinfomgt.2020.102242

j. **Liu, B.** (2012). "Sentiment Analysis and Opinion Mining." *Synthesis Lectures on Human Language Technologies*, 5(1), 1-167. doi: 10.2200/S00416ED1V01Y201204HLT016

k. **Riley, S., & Finkelstein, A.** (2020). "Modeling the Impact of Public Health Campaigns on Social Media Use During COVID-19." *Journal of Health*

Communication, 25(10), 766-773. doi: 10.1080/10810730.2020.1820155

l. **Barbieri, F., & Saggion, H.** (2020). "COVID-19 Tweets: Automatic Identification of Online Misinformation." *Proceedings of the International Conference on Artificial Intelligence*, 243-249. Retrieved from https://arxiv.org/abs/2004.04691

m. **Poria, S., & Cambria, E.** (2016). "Sentiment Analysis and Opinion Mining." *Springer International Publishing*. ISBN 978-3-319-44463-0.

n. **Abadi, M., & Others**. (2016). "TensorFlow: A System for Large-Scale Machine Learning." *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation*, 265-283. Retrieved from https://www.usenix.org/system/files/conference/osdi16/osdi16-abadi.pdf

o. **Xu, W., & Kim, Y.** (2020). "Emerging Trends in the COVID-19 Crisis: Text Mining and Social Media Sentiment Analysis." *Journal of Applied Computing and Informatics*, 19(4), 102-115. doi: 10.1016/j.jaci.2020.06.003

p. **Google Fact Check Tools**. (2021). "Fact-Checking for a Safer Online Community." Retrieved from https://toolbox.google.com/factcheck/explorer

q. **Pushshift API Documentation**. (2021). "Pushshift Reddit Data API." Retrieved from https://pushshift.io

r. **Apache Kafka Documentation**. (2020). "Apache Kafka: A Distributed Streaming Platform." Retrieved from https://kafka.apache.org/documentation/

s. **Rohde, D. L. T., & Bi, X.** (2020). "Predicting Health Trends from Social Media: COVID-19 and Beyond." *Journal of Healthcare Informatics Research*, 4(2), 137-148. doi: 10.1007/s41666-020-00051-1

t. **RoBERTa (Robustly Optimized BERT Pretraining Approach)**. (2019). "RoBERTa: A Robustly Optimized BERT Pretraining Approach." *arXiv preprint arXiv:1907.11692*. Retrieved from https://arxiv.org/abs/1907.11692