Analysing the impact of Machine Learning Health Operations (MLHOps) on Mental Health and Stress

MSc Research Project
Cloud Computing

Naaga Barani Govindan Venkatraj
Student ID: x22104038

School of Computing
National College of Ireland

Supervisor: Ahmed Makki

National College of Ireland
Project Submission Sheet
School of Computing

Student Name:	Naaga Barani Govindan Venkatraj
Student ID:	x22104038
Programme:	Cloud Computing
Year:	2023
Module:	MSc Research Project
Supervisor:	Ahmed Makki
Submission Due Date:	14/12/2023
Project Title:	Analysing the impact of Machine Learning Health Operations (MLHOps) on Mental Health and Stress
Word Count:	7725
Page Count:	22

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.

ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature:	Naaga Barani Govindan Venkatraj
Date:	13th December 2023

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST:

Attach a completed copy of this sheet to each project (including multiple copies).	☐
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	☐
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	☐

Assignments that are submitted to the Programme Coordinator office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Analysing the impact of Machine Learning Health Operations (MLHOps) on Mental Health and Stress

Naaga Barani Govindan Venkatraj
X22104038

Abstract

This study examines how Machine Learning Health Operations (MLHOps) can be used to address the growing global mental health issues, which have been made worse by the COVID-19 pandemic. Acknowledging mental health as integral to overall well-being, the study emphasizes the urgent need for innovative solutions and early intervention. Leveraging the power of cloud computing, specifically Microsoft Azure, the research aims to utilize social media data, particularly Twitter, for early detection of mental health conditions, with a focus on depression. The dataset, collected from multiple Twitter accounts, was comprised of normal tweets and a labelled dataset of depressive tweets from users diagnosed with depression and other mental health conditions. Natural language processing (NLP) techniques were applied to extract features from the tweet text that could indicate signs of depression. The tweets were preprocessed by removing URLs, usernames, hashtags, and stopwords, then vectorized into dense word embeddings using the Spacy language model. Three supervised machine learning models were trained on the dataset to classify users as depressed or not depressed based on these text features - K-Nearest Neighbors (KNN), Random Forest (RF), and Support Vector Machine (SVM). The system achieved an accuracy of 82% with RF, outperforming KNN at 78% and SVM at 85%. To operationalize the analysis, an end-to-end Azure Machine Learning (AML) pipeline was built comprising data ingestion, preprocessing, model training and evaluation to reflect the ML Health Operations (MLHOps) practices. This enables an automated workflow to continuously collect new Twitter data, extract features, and identify users at risk of depression. The methodology demonstrates the feasibility of monitoring population-level mental health from social media data. Future work includes expanding beyond depression to detect other conditions and improving model feasibility.

1 Introduction

Mental wellness is a fundamental aspect of our psychological health. The World Health Organization (WHO) defines mental health as the condition of being in a state of well-being where individuals can recognize and utilize their capabilities, effectively manage typical daily pressures, perform productively, and make positive contributions to their communities (WHO. 2020). In 2018, the prevalence of depression reached over 300 million individuals worldwide, making it the primary cause of global disability. Depression incurs various repercussions, encompassing both individual and societal expenses. Depression can sometimes result in thoughts of suicide and actual attempts to end one's life (Vigo et al.; 2016).

Social networking has emerged as a significant reservoir of medical data, enabling us to identify and anticipate affective disorders, which might serve as a supplementary instrument to monitor mental health and conduct information surveillance. The development of Machine Learning (ML) and Artificial Intelligence (AI) has radically altered several industries, including the healthcare industry. Machine Learning Health Operations (MLHOps) is an emerging approach that amalgamates machine learning with the operational aspects of healthcare to enhance healthcare delivery, streamline operations, and ultimately improve patient outcomes. The application of MLHOps in the domain of mental health has the potential to be a game-changer.

Moreover, the utilization of diverse approaches grounded regarding the topics of NLP and ML model frameworks have demonstrated their efficacy in providing support and assistance by automating the process of detecting initial indicators of mental illness by technological means by examining the content posted on the internet by individuals (Conway and Conner; 2016). Twitter is a highly significant social media network in terms of its user base, boasting over 330 million active users globally. Starting from November 2017, the maximum character limit for a tweet has been raised from 140 to 280. Through the examination of vast quantities of text, scholars can establish a connection between the utilization of language in daily life and social conduct and individuality. Hence, examining the content shared on social media sites might yield insights into various personality traits, lifestyles, and psychological problems (Valdez et al.; 2020). However, individuals with depression may also experience discomfort when engaging in social interactions and consuming content on social media platforms (Prieto et al.; 2014). The detection and monitoring of mental diseases, such as depression, can be facilitated by analysing many aspects of the messages, including the quantity and frequency of tweets, their distribution throughout the day and night, and their seasonal patterns (De Choudhury.; 2013). This knowledge can assist healthcare practitioners and health institutions in making informed decisions to enhance the management of people with depression.

This research aims to develop a system to identify Twitter users with depression based on analysis of tweet text. Tweets will be collected using the Twitter API and Tweepy library. A labelled dataset with tweets from users diagnosed with depression is compiled from Kaggle. Natural language processing is applied to extract features from tweets which are then used to train machine learning models - KNN, RF, and SVM. The models will be evaluated and optimized to maximize accuracy, F1-score and other metrics in classifying users as depressed or not depressed. To operationalize the analysis at scale, an end-to-end Azure Machine Learning pipeline will be implemented for automated data ingestion, preprocessing, model training/retraining, and deployment to deploy the MLOps deployment model. This will enable real-time analysis of streaming Twitter data to identify users at risk of depression.

1.1 Motivation

The motivation for embarking on this research study is rooted in a complex interplay of critical factors that collectively underscore the pressing need to address mental health challenges and leverage innovative technologies for the betterment of healthcare. With a significant percentage of the world's population suffering from disorders like depression and anxiety, mental health concerns have become an international epidemic. Early intervention in the identification of individuals at risk of mental health conditions is essential, as it holds the potential to prevent the exacerbation of symptoms and alleviate suffering. In a period when technology is changing every aspect of every aspect of our lives, cloud computing and machine learning health operations (MLHOps) play an essential part in the healthcare sector. Furthermore, the COVID-19 pandemic has made the epidemic of mental illness worse, calling for a targeted intervention. Leveraging the vast data generated on social media platforms, such as Twitter, for early mental health detection is a promising avenue. Microsoft Azure, a leading cloud computing platform, offers a robust suite of tools and services that align with the research's requirements for scalability and cost-efficiency. By bridging the technological advantages of MLHOps with Azure (Pavlova and Berkers; 2020). This motivates developing systems to analyze social media data to identify signs of emerging mental health issues.

1.2 Research Question

Q: How does the integration of Machine Learning Health Operations (MLHOps) on the Microsoft Azure cloud platform impact the early detection and assessment of mental health conditions, particularly depression, through the analysis of Twitter data, and what are the benefits of Azure in continuous integration and monitoring for this application?

1.3 Research Objective

Objective	Description
Objective 1	Develop a comprehensive framework that integrates MLHOps with Azure cloud services for mental health assessment using Twitter data.
Objective 2	Identifying signs of depression by developing and evaluating NLP and ML-based frameworks to accurately identify signs of depression from Twitter data, specifically focusing on English language tweets.
Objective 3	Document and present the practical benefits of integrating Azure cloud services and MLHOps in terms of improving mental health support

1.4 Cloud Computing and Azure

To accomplish the objectives of this research, a robust and scalable computing infrastructure is essential. Cloud computing has emerged as the preferred platform for processing data, storing it, and hosting machine learning models. Microsoft Azure, a leading cloud computing platform, offers a full range of resources and services to make the development of machine learning and projects regarding analytics using data easier. Cloud computing has heralded a paradigm shift in the way businesses operate and manage their IT infrastructure. It has evolved from a simple advancement in technology to an important strategic resource that can boost productivity, adaptability, and cost reduction across industries throughout the past decade. This transformation is largely attributed to the advantages offered by cloud computing. One of the primary advantages of cloud computing is its scalability. Organizations can easily scale up or down their resources to meet changing demands, ensuring that they pay only for what they use.

This flexibility is invaluable in research and healthcare, where computational requirements can fluctuate significantly based on the scope of the study and the amount of data to be processed. Machine Learning Health Operations (MLHOps) is an innovative framework that combines machine learning models with operational practices. It streamlines the deployment, monitoring, and management of machine learning models, making it particularly well-suited for healthcare applications. This research will explore how MLHOps can be harnessed to improve mental health assessment and intervention processes.

1.5 Thesis structure

The structure of this research report has six distinct sections:

Section 1: Introduction

- Background & Motivation

Section 2: Research Question

- Objectives

Section 3: Literature Review

- Limitation & Gaps

Section 4: Methodology

- Design & Implementation

Section 5: Evaluation

- Result & Discussion

Section 6: Conclusion

- Summary & Future works

2 Related Work

2.1 Social Media and Mental Health

Karim et al. (2020) evaluated sixteen cross-sectional, long-term, qualitative in nature, and systematic studies on the relationship between mental health and social media. Time spent on social media may have both positive and negative consequences on mental health, notably anxiety and sadness. The report suggested more qualitative research and vertical cohort studies to better understand social media's effects on mental health and improve mental health solutions. A correlational study (Berryman et al.; 2018) examined social media's potentially harmful effects on young people's mental health. Except for vague booking, which predicted suicidal ideation, social media use did not predict mental health problems. The findings suggest that fear about social media use might not be justified—apart from ambiguous booking, which could point to severe mental health issues.

Gao et al. (2020) examined mental health issues and social media use in China during the COVID-19 pandemic. Chinese adults aged 18 and older had high rates of sadness, anxiety, and CDA throughout the outbreak, according to the study. More than 80% of people said they used social media frequently, and frequent use was linked to greater risks of anxiety and CDA than infrequent use. According to the research, authorities should give depressive disorders and anxiety in people of all ages priority during medical emergencies. Recent research by (Conway and Conner. 2016) used social media big data, NLP, and ML for population-level mental health surveillance and research. They advised using social media analysis to derive population-level inferences from naturalistic first-person reports of user behaviour and opinions that may indicate mental health state. However, using social media for health research presents ethical concerns about user privacy and public/private content. An eight-year follow-up investigation examined teenage online social networking use and mental wellness (Coyne et al.; 2020). It used cross-lagged pathways, within-time covariances, latent covariances, and latent means to analyze how social media use affects depression and anxiety. Social media use was linked to mental health markers, using parameter estimates and fit indices. The study's findings help explain how social media use may affect adolescents' mental health and the use of social media datasets for analysis of mental health and stress.

2.2 Twitter Datasets in Mental Health Analysis

Rahman et al. (2022) proposed a study that utilized sentiment analysis, ML techniques, and emotional classification to analyse the tweets of four mental health-related NGOs in Malaysia. The study used WordCloud techniques to visualize the most frequent positive and negative words used by the tested NGO Twitter accounts, providing a visual representation of the sentiment words. According to the study's outcomes, NGOs should use more uplifting text posts on social media to teach people about mental health problems and raise awareness of these issues. (Naseem et al.; 2021) developed a large-scale manually annotated COVID sentiment data set, demonstrating how public sentiments concerning Coronavirus were traced, extracted indicative topics, and benchmarked the performance of different state-of-the-art ML text classification mechanisms. The key results included the development of the COVIDSenti data set consisting of 90,000 tweets labelled into positive, negative, or neutral sentiment classes, the demonstration of how public sentiments concerning Coronavirus were traced, the extraction of indicative topics, and the benchmarking of the performance of different state-of-the-art ML text classification mechanisms. Additionally, the study performed feature extraction using vectorization techniques and word embeddings, and it employed machine learning and deep learning-based classifiers to gauge performance in the sentiment classification task.
Levanti et al. (2023) examined the frequency of depressive and anxiety disorders on Twitter throughout the COVID-19 stay-at-home duration in seven major U.S. cities. The study used AI-based language assessments to analyze the tweets and compared the results between 2019 and 2020. The authors discovered that symptoms of depression and anxiety were higher in 2020 than they were in 2019, which may indicate that the worldwide epidemic affected mental health. The study provides city-level word, phrase, and topic scores for each month, as well as standardized monthly scores for depression and anxiety for each city. The research is based on a large sample of Twitter users and tweets, and the findings are relevant for understanding the mental health implications of the COVID-19 pandemic. The study conducted by (Lauricella et al.; 2022) collected and analyzed tweets from February 2019 to March 2021, encompassing the period before and during the pandemic. To develop a fresh theoretical perspective on the mental health concerns of students, the investigation used a solid theoretical method. The study suggests that systemic issues in higher education directly impact student wellness and proposes

reevaluating structures such as grades, community, collaboration, and assessment to address stress and anxiety. The study's limitations, such as the utilization of Twitter data and the platform's public nature, were also addressed. The report further highlighted the imperative of addressing systemic factors that contribute to students' mental health and the possibility of implementing transformative modifications in the higher education journey.

2.3 ML and NLP in Mental Health Analysis

Rahman et al. (2020) presented a critical analysis of mental health detection in Online Social Networks (OSNs) based on data sources, ML, and feature extraction methods. The authors emphasize the potential of OSNs as a data source for early detection of mental health problems and the need for comprehensive adoption, innovative algorithms, and computational linguistics to overcome the challenges associated with mental health detection. It provides insights into the use of feature extraction methods such as Term Frequency-Inverse Document Frequency N-Gram, Bag of Word (BoW), Word2Vec, Global Vectors for Word Representation (GloVe), Linguistic Inquiry and Word Count (LIWC), and Latent Dirichlet Allocation (LDA) in text classification for mental health problem detection. (Le Glaz et al.; 2021) discussed the application of ML and NLP methods in mental health research and clinical practice. It analyzed a total of 58 studies that utilized ML and NLP techniques for mental health purposes. The investigations encompassed a range of subjects, including the extraction of symptoms, the classification of sickness severity, the comparison of therapy effectiveness, and the provision of psychopathological clues. The review also examined the potential of NLP and ML techniques in mental health research and therapeutic practice, emphasizing their utility and the ethical considerations that must be taken into account. (Govindasamy and Palanichamy. 2021) presented a study aimed to detect depression in users based on their Twitter data using ML algorithms, specifically Naïve Bayes and a hybrid model, NBTree. The results of the study show that both algorithms perform equally, demonstrating the potential of using machine learning for depression detection on social media. A thorough description of depression, the use of machine learning to identify depression through social media, and the study methodology which included data collection, preprocessing, sentiment analysis, and classification were all covered in various points in the report.

Kumar et al. (2022) employed various features such as sentiment analysis, emoticon usage, and linguistic features to identify depressive symptoms in user posts. They proposed a numerical score for each user based on the sentiment value of their tweets and demonstrated that this feature can detect depression with an accuracy of 78% with the XGBoost classifier. To achieve excellent accuracy in depression detection, the study's main contributions include reviewing the literature on various emotion detection methods, choosing linguistic, subject matter, emoticon, and sentiment characteristics for the research problem, defining an indicator to assign a numerical grade to each user based on the sentiment value of their tweets, and using machine learning techniques to demonstrate the benefits of combining various characteristics. (Prakash et al.; 2021) discussed the use of ML algorithms, specifically SVM and RF to analyze tweets for self-assessed depressive features, aiming to aid in the early detection and treatment of depression. The authors collected real-time tweets, preprocessed the data, and classified the tweets using the mentioned algorithms. The proposed model works by synchronizing different machine learning algorithms to work as an ensemble model for higher efficiency and accuracy. (Owen et al.; 2020) presented the study on preemptive detection of depression and anxiety in Twitter by applying state-of-the-art classification models to this dataset, providing a competitive set of baselines alongside qualitative error analysis. They evaluated several binary classifiers on the datasets, including a SVM with TF-IDF features and a classifier based on the average of word embeddings within the tweet. They also used pre-trained language models (LMs) like BERT and ALBERT. The results showed that the SVM with word embeddings achieved an accuracy of 72.7%, outperforming the SVM with TF-IDF features, which achieved an accuracy of 63.3%.

Perez et al. (2022) presented a novel approach for automatically estimating the degree of depression in social media users using neural language models and word embeddings. The first approach takes the use of the generic language that users have used in their writings, whereas the second looks for more concrete proof from publications that address the symptoms in question. The Beck Depression Inventory (BDI-II) total score is automatically estimated using both approaches. Using benchmark Reddit data, the scientists discovered that neural language model-based techniques provide a workable substitute for predicting depression rating scales, even in situations when there is a dearth of training data. (Tejaswini et al.; 2022) discussed the increasing prevalence of depression and the challenges in its early detection. It focuses on the use of NLP techniques and deep learning (DL) models to analyze text content from social media for the early detection of depression. To overcome the difficulties of accurately identifying depression from literature, researchers presented a novel hybrid deep learning neural network architecture dubbed "Fasttext Convolution Neural Network with Long Short-Term Memory (FCL)". The study article makes a significant addition to the field of mental health studies and technological advances as it focuses on the use of NLP and deep learning techniques for early identification of depression from social media content. (Chakraborty et al.; 2020) discussed the effects of COVID-19 on both physical and mental well-being throughout the world, the role that social media plays in data dissemination, and the necessity of relevant and reliable material during the epidemic. It presents an analysis of tweets related to COVID- 19, highlighting the prevalence of negative sentiments in retweeted tweets and the ineffectiveness of social media in guiding people during the pandemic. The paper also proposes a model using deep learning classifiers and a Gaussian membership function based fuzzy rule base to identify sentiments from tweets, with a focus on the need for fact-checking and the responsible sharing of information on social media. (Tadesse et al.; 2019) addressed the early detection of suicide ideation through deep learning and machine learning-based classification approaches applied to Reddit social media. They employ an LSTM-CNN combined model to evaluate and compared to other classification models, showing that the combined neural network architecture with word embedding techniques can achieve the best relevance classification results. The main contributions of this research are N-gram analysis, classical features analysis, and comparative evaluation of the proposed model with other DL and traditional ML classifiers.

2.4 MLOps framework

Kreuzberger et al. (2023) discussed the concept of Machine Learning Operations (MLOps) and its application in industrial machine learning (ML) projects. The paper addresses the challenges of automating and operationalizing ML products and aims to provide a comprehensive definition of MLOps, along with the necessary principles, components, roles, architecture, and workflows. The paper also drew parallels between MLOps and the DevOps paradigm, emphasizing the importance of automation, collaboration, and continuous monitoring in MLOps. The paper provides an in-depth analysis of these principles and their implementation within the technical components of MLOps. (Garg et al.; 2021) presented the integration of Continuous Integration and Continuous Delivery (CI/CD) practices into MLOps, the differences between DevOps and MLOps, and the use of GitOps for model deployment. The paper also presents the concept of MLOps, its levels, and open research challenges. It provides an overview of MLOps platforms such as KubeFlow, Amazon SageMaker, and MLFlow, and discusses the importance of monitoring AI models and the challenges in MLOps implementation. (Chowdary et al.; 2022) observed the challenges in automating and operationalizing ML products and introduces the paradigm of MLOPs to address these issues. It includes an overview of the necessary principles, components, roles, architecture, and workflows of MLOps, as well as a comprehensive definition and open challenges in the field. The authors conducted mixed-method research, including a literature review, a tool review, and expert interviews, to provide clear guidelines for professionals and researchers in the field of MLOps.

Qasem et al. (2015) presented the application of sentiment classification of Twitter data to predict stock market variables. It evaluates the general accuracy of two machine learning techniques (neural networks and logistic regression) in classifying tweets about stocks as neutral, negative, or positive. The study uses a dataset of 42,000 automatically annotated tweets covering four technology-related stocks collected from Twitter. The classifiers achieved the same overall accuracy (58%), but empirical experiments showed that using Unigram TF-IDF outperforms TF. The whole work was implemented as an ML pipeline in Microsoft Azure and discusses the deployment in the context of the research. (Chaudhary et al.; 2023) proposed the use of Twitter data to characterize the pandemic-related experiences of the United States' African American population using aspect-based sentiment analysis. The study leverages a machine learning pipeline to filter and analyze nearly 4 million tweets, revealing that most tweets had a negative tone. The effect of the coronavirus pandemic on African Americans is further discussed, with a focus on the pressing need to investigate their perceptions, actions, and attitudes about COVID-19.

To summarize the literature study, a significant amount of research has been done over the years to examine depression and enhance mental health in people using the vast resource pool of social media data, particularly Twitter.. Both ML and NLP approaches have been proposed and implemented by various authors for diverse applications to detect and classify the gathered social media data. Recently, there has been an influx of large language models (LLM) that can analyze Twitter data faster and more accurately, but they are more computation-intensive and complex in their architecture. Hence, this work was proposed in the light of using supervised

ML models like RF, KNN and SVM and word embedding models like SpaCy to train the processed Twitter data corpus and identify the signs of depression and poor mental health. However, as presented in the previous section, this work also attempts to implement this depression detection model as an ML pipeline in an MLOPs deployment framework (Azure MLOPs) to facilitate the accelerated development of the model and future large-scale deployments.

3 Research Methodology

The methodology section covers the key aspects of data collection, preprocessing, model development and deployment. Figure.1 provides a model workflow diagram to help visualize the steps involved in the design technique.

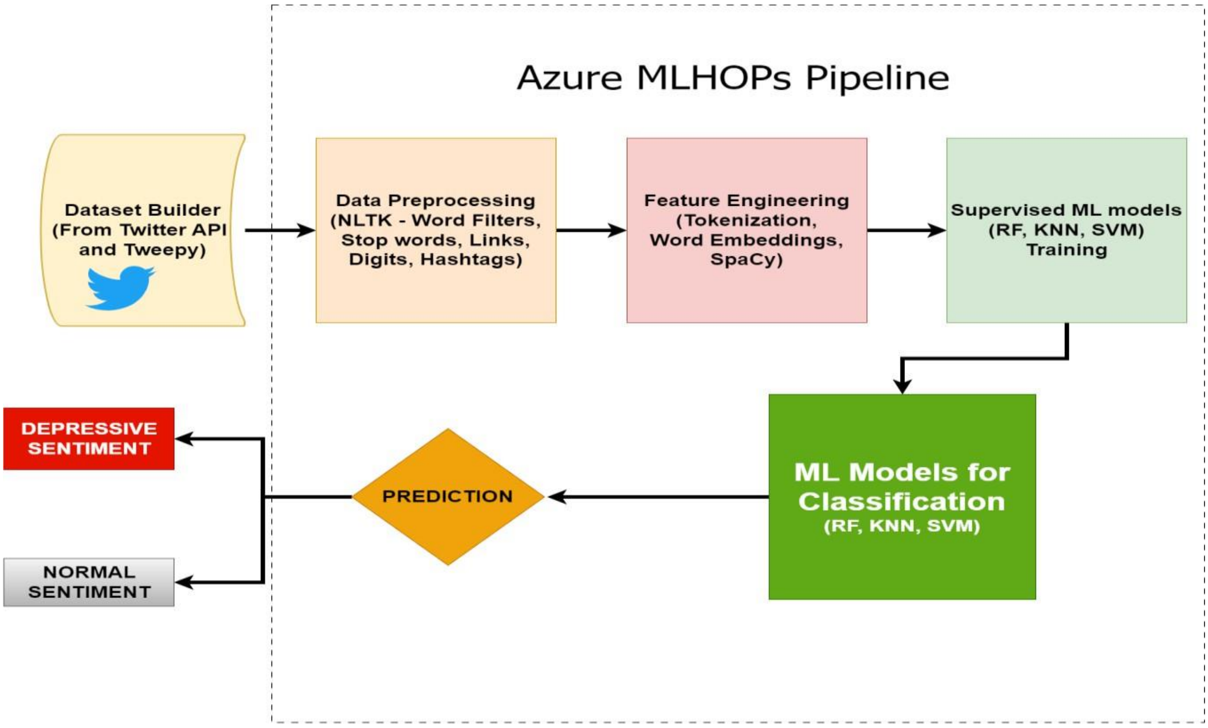

Figure 1. Research Methodology Workflow

3.1 Data Collection:

The data collection process involves gathering the tweets from the Twitter API and Tweepy libraries in addition to the Kaggle dataset. The use the Twitter API v2 to collect tweets on topics of interest where we need to gather both normal and depressive tweets to provide a balanced dataset to the model. This includes establishing a project, granting developer access, searching for recent or old tweets using APIs, and storing the JSON result as a CSV file.

Alternatively, an existing Twitter dataset from Kaggle, such as the sentiment analysis dataset[1]. This dataset is also used in addition to the prepared ones to provide more randomness to the tweets with different sentiment expressions.

3.2 Data Preprocessing:

The next step is to preprocess the collected tweets to remove any irrelevant information, such as retweets, links, and non-English tweets. This step also includes text normalization, tokenization, and removing stop words. NLTK (Natural Language Toolkit) python library allows us to tokenize tweet messages, which processes Twitter handles, phone numbers, case in sentences, and emoticons. Later, text cleaning is performed which includes tasks such as removing URLs, mentions, digits, stop words, and punctuation. Hashtags are extracted from tweets, which can provide valuable information for sentiment analysis and topic modelling.

3.3 NLP for feature extraction

spaCy, an NLP library is used for information extraction, natural language understanding systems, and preprocessing text for ML classification. It offers various linguistic features, including Named Entity Recognition (NER), Part-of-Speech (POS) tagging, dependency parsing, word vectors etc., and uses the vector representations of words for feature extraction. The library provides access to pre-trained word vectors, such as the "en_core_web_lg" model, which includes 300-dimensional word embeddings for over 1 million unique tokens. These word vectors can be used for various NLP tasks, such as text classification, clustering, and similarity analysis. Additionally, SpaCy's word embeddings can be used for tasks like sentence classification by obtaining the mean vector for an entire sentence, which is useful for capturing the overall semantic meaning of the sentence. It provides a powerful tool for representing words as dense vectors, enabling the application of advanced NLP techniques.

3.4 Machine Learning for Classification

This is focused on comparing the performance of different supervised ML techniques, including KNN, RF and SVM for classifying tweets as depressive or not. These algorithms have been extensively used and validated for tweet classification tasks in various studies and demonstrate their suitability for handling tweet data as discussed in section 2.3. All three models have hyperparameters that can be tuned for improved performance. This makes them adaptable to the tweet classification task. They can handle multi-class classification problems well, which is important as tweets need to be classified into several possible categories. The models are evaluated in terms of accuracy, precision, recall and F1-score for performance evaluation.

3.5 MLHOps using Azure ML Pipelines

This research leverages Azure ML pipelines to create reproducible and reusable workflows for model training, evaluation, and deployment. The pipelines are created for the different stages

[1] https://www.kaggle.com/datasets/ywang311/twitter-sentiment/data

of the model development that involves data preprocessing, model training and prediction. The idea is to deploy the trained classification model using Azure ML, which provides state-of-the- art ML pipelines for consistent model delivery with increased fault tolerance. The below diagram illustrates the workflow of MLOps in Mental Health Analysis.

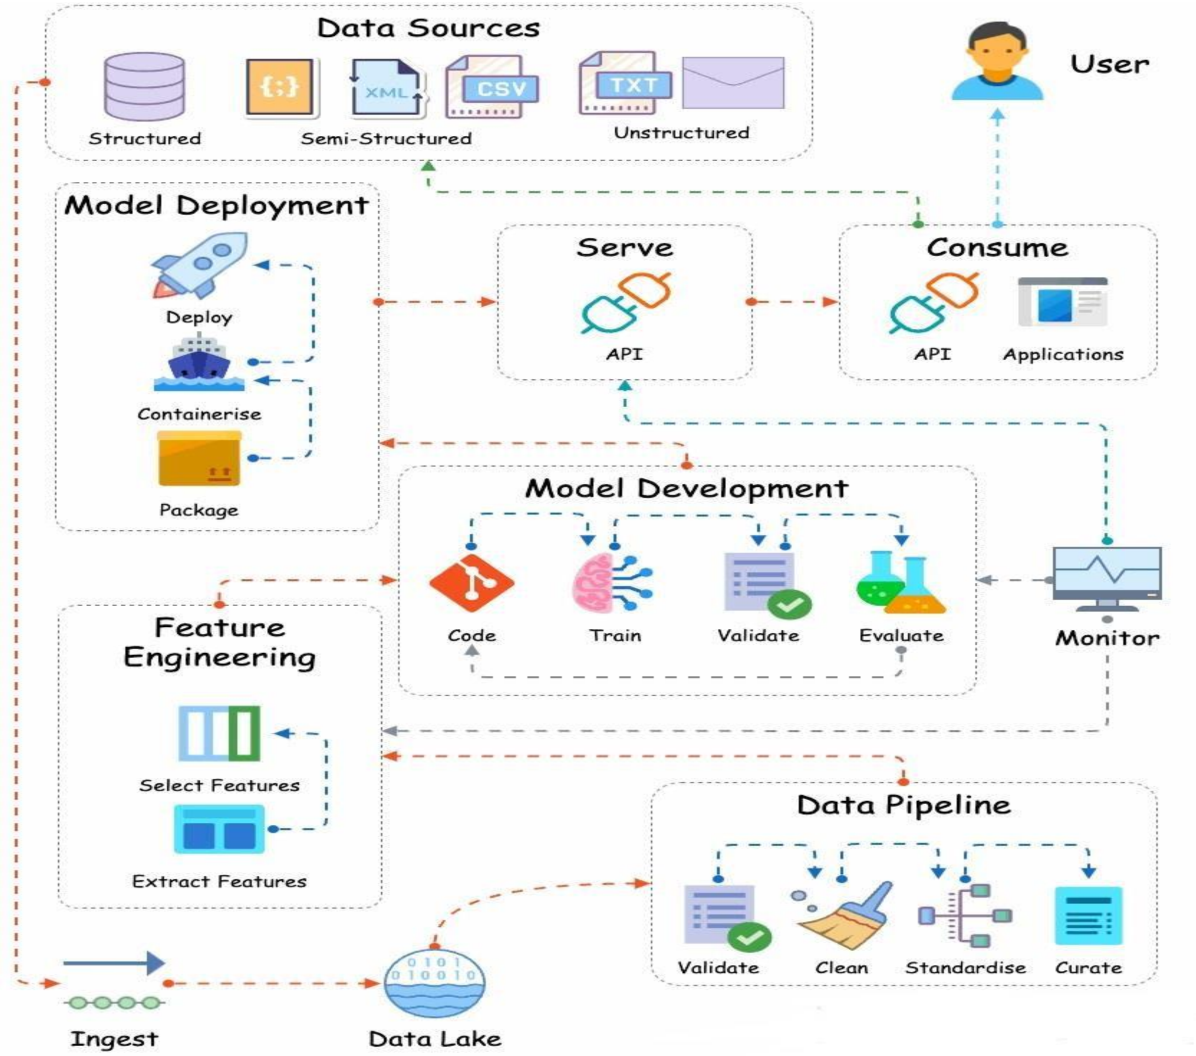

4 Model Specifications

4.1 Random Forest (RF)

Random Forest is an ensemble learning method that operates by constructing a multitude of decision trees at training time. For classification tasks, the output of the random forest is the class selected by most trees, while for regression tasks, the mean or average prediction of the individual trees are returned. To produce a forest with no correlation of decision trees, it extends the bagging technique by combining feature randomness with bagging.

Feature randomness, also known as feature bagging or the random subspace method, generates a random subset of features, ensuring low correlation among decision trees. This is a key difference between decision trees and random forests. Random forests are commonly used for both classification and regression problems and are known for their ease of use and flexibility. Figure 2 below presents the structure of the random forest model.

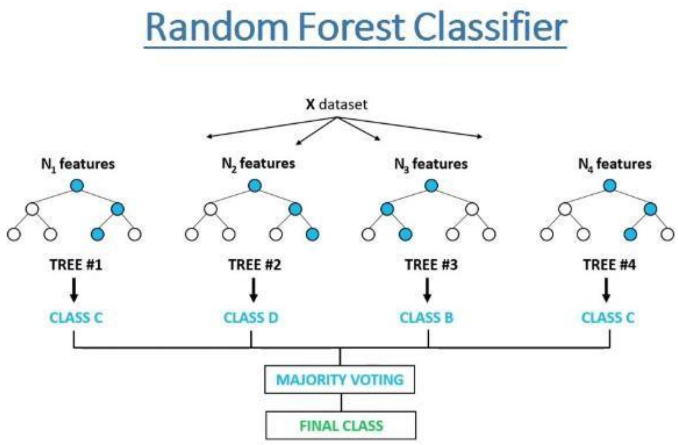

Random Forest has been effectively used for tweet classification in several studies For example, a study on the categorization of catastrophe tweets using Random Forest showed that it is useful for precisely determining the location of the disaster and the areas that need assistance in the disaster zone. (Kanimozhi and Sara; 2022). Additionally, the Random Forest Twitter Bot Classifier proposed a set of attributes for a Random Forest classifier that resulted in high accuracy (90.25%) and generalizability (Schnebly and Sengupta; 2019). These findings highlight the suitability of Random Forest for tweet classification tasks, particularly in scenarios such as disaster management and sentiment analysis.

4.2 Support Vector Machines (SVM)

SVM is a supervised ML algorithm used for both classification and regression tasks. For the technique to function, data points are plotted as points in a space with n dimensions, with each feature's value representing a specific coordinate. The classification into respective categories is done by finding the optimal hyperplane that differentiates the two classes in the best possible manner. The algorithm aims to find a boundary that maximizes the margin between the two classes. Support vectors are the elements of the training set that would change the position of the dividing hyperplane if removed. They are the essential components of the training set that directly affect where the decision layer should be placed.

SVM has been effectively used for tweet classification in several studies. (AminiMotlagh et al.; 2022 examined an ensemble system of classification that predicted tweet sentiment using SVM in addition to other models like Random Forest and Naive Bayes. SVM performed well in English and Spanish datasets. SVM has proven to be an effective model for handling noise and complexity in tweet classification tasks like sentiment analysis, topic classification etc. It works well with text data since it can maximize margins.

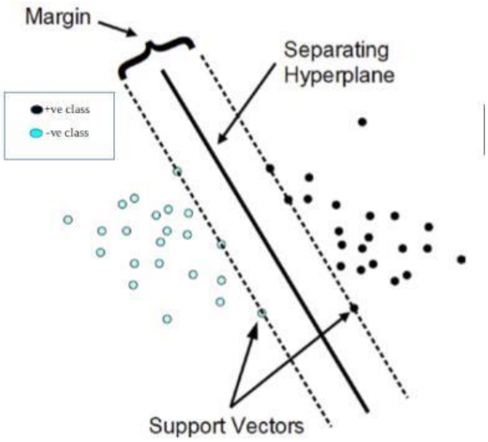

Figure 3. SVM Classifier

4.3 K-Nearest Neighbours (KNN)

The KNN algorithm is a supervised machine-learning approach used for classification and regression tasks. It operates by finding the distance between the mathematical values of data points. The algorithm stores all available cases (test data) and classifies new cases based on a similarity measure. When implementing KNN, the first step is to transform data points into their mathematical values (vectors). To determine the likelihood that two data points would be similar, the algorithm first calculates the distance that exists between every point of data and the test data. The algorithm assumes that all data points are geometrically close to each other, and it classifies new points based on the classes of their nearest neighbours. The majority vote of the K nearest neighbours is used to classify new points.

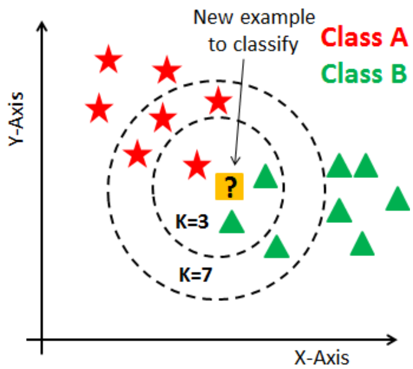

Figure 4. KNN Classifier

KNN has been utilized for sentiment analysis and tweet classification in various studies. a study applied KNN for sentiment analysis of Twitter data related to US airlines, achieving promising results in classifying tweets as positive, negative, or neutral[2].

[2] https://towardsdatascience.com/sentiment-analysis-of-twitters-us-airlines-data-using-knn-classification-91c7da987e13

4.4 NLTK Library

NLTK[3] is an open-source Python library for NLP tasks with interfaces to text processing libraries and lexical resources like WordNet to support NLP tasks like classification, tokenization, stemming, tagging, parsing and semantic reasoning. It has graphical demonstrations and sample data sets to explain the concepts and applications of NLP. NLTK supports key NLP tasks like part-of-speech tagging, named entity recognition, building tree models for text structure analysis, sentiment analysis, and vector space modelling. NLTK provides a practical platform for applying NLP using Python as shown below.

```
import nltk
```

4.5 spaCy Library

The spaCy library is a popular open-source Python library for NLP tasks. It comes with built-in word embedding models, which are pre-trained vector representations of words. These word embeddings can be used for various NLP tasks, such as text classification, clustering, and similarity analysis. To use spaCy's word embeddings, a pretrained model that includes word vectors, such *as en_core_web_md or en_core_web_lg* can be loaded into the python code.

```
import spacy

# Load the spaCy model
nlp = spacy.load("en_core_web_md")

# Process a sentence using the model
doc = nlp("Sample Processing with Spacy")

# Get the vector for a specific word
word_vector = doc[3].vector

# Get the mean vector for the entire sentence
sentence_vector = doc.vector
```

In the above code script, the *nlp* object is used to process a sentence, and the word vectors can be accessed through the *vector* attribute of the tokens in the *doc* object. The vector representation attribute of the *doc* item itself may be used to get the mean *vector* for the full sentence.

5 Implementation

The project implementation phase involves developing and deploying the ML models that can effectively detect tweets containing depressive characteristics using python ML libraries

[3] https://www.nltk.org/install.html

including *spaCy*, *NLTK*, *Scikit-learn* and other data handling and processing libraries like *pandas*, *NumPy and MatplotLib*.

5.1 Data preprocessing:

```
] depressive_tweets_df = pd.read_csv("/content/drive/MyDrive/MLHops/Depression_Tweets/data/depressive_tweets.csv")
  random_tweets_df = pd.read_csv("/content/drive/MyDrive/MLHops/Depression_Tweets/data/random_tweets.csv",
                                 encoding = "ISO-8859-1", usecols = range(0,4), nrows = 40000)
```

```
] depressive_tweets_df
```

	Unnamed: 0	tweet.id	created_at	text	location	retweet	favorite
0	0	1447537898572574730	2021-10-11 12:21:43	Open discussion. Between the Transfer Portal a...	Cheyenne Wyoming	0	0
1	1	1447540582490988553	2021-10-11 12:32:23	Plenty of things are changing in my life and t...	NaN	0	0
2	2	1447807717859491842	2021-10-12 06:13:53	I feel a little hopeless. Anyone else? #hopele...	NaN	0	0
3	3	1448076026219692033	2021-10-13 00:00:03	Which is more healthy? Hope, or hopelessness? ...	Denver, CO	0	0
4	4	1448382047375040513	2021-10-13 20:16:04	So someone tell me how do I get over #HOPELESS...	Portland Or .	0	2

Figure 5. Load Dataset

```
[ ] ## Drop unnecessary columns
    depressive_tweets_df.drop(columns=['Unnamed: 0'], inplace=True)
    new_rand_df.drop(columns=['ï»¿ItemID', 'index','Sentiment', 'SentimentSource'], inplace=True)
```

Figure 6. Drop Unwanted Columns

```
## Finding unique values in each column
for col in depressive_tweets_df:
    print("There are ", len(depressive_tweets_df[col].unique()), "unique values in ", col)
```

```
There are  18190 unique values in  tweet.id
There are  18071 unique values in  created_at
There are  17107 unique values in  text
There are  4648 unique values in  location
There are  74 unique values in  retweet
There are  159 unique values in  favorite
```

Figure 7. Finding unique values in each column

	text	label
0	Open discussion. Between the Transfer Portal a...	1
1	Plenty of things are changing in my life and t...	1
2	I feel a little hopeless. Anyone else? #hopele...	1
3	Which is more healthy? Hope, or hopelessness? ...	1
4	So someone tell me how do I get over #HOPELESS...	1

Figure 8. Cleaned Dataset

The *tweets_cleaner()* function (Figure 9) is used to perform the stepwise cleaning process that includes the following:

- Change all letters to lowercase.
- If url links, then don't append to the dataset.
- remove hashtag, @mention, emoji, and image URLs.

- remove punctuation.
- stopwords for *english* and create word tokens.

```
## Function to perform stepwise cleaning process
def tweets_cleaner(tweets):
  cleaned_tweets = []
  for tweet in tweets:
    tweet = tweet.lower() #lowercase

    # if url links then don't append to avoid news articles
    # also check tweet length, save those > 5
    if re.match("(\w+:\/\/\S+)", tweet) == None and len(tweet) > 5:

      #remove hashtag, @mention, emoji and image URLs
      tweet = ' '.join(re.sub("(@[A-Za-z0-9]+)|(\#[A-Za-z0-9]+)|(<Emoji:.*>)|(pic\.twitter\.com\/.*)", " ", tweet).split())

      #fix weirdly encoded texts
      tweet = ftfy.fix_text(tweet)

      #expand contraction
      tweet = expandContractions(tweet)


      #remove punctuation
      tweet = ' '.join(re.sub("([^0-9A-Za-z \t])", " ", tweet).split())

      #stop words and lemmatization
      stop_words = set(stopwords.words('english'))
      word_tokens = nltk.word_tokenize(tweet)

      lemmatizer=WordNetLemmatizer()
      filtered_sentence = [lemmatizer.lemmatize(word) for word in word_tokens if not word in stop_words]
      # back to string from list
      tweet = ' '.join(filtered_sentence) # join words with a space in between them

      cleaned_tweets.append(tweet)
```

Figure 9. Tweet Cleaner Function

The cleaned dataset is then saved to "*preprocessed_data.csv*" file for further use during the model training phase.

5.2 Model Training:

The model training phase begins with downloading the word embedding model from spacy using, *python -m spacy download en_core_web_lg*.

```
## libraraies for classification
from sklearn.pipeline import Pipeline
import sklearn.metrics as skm
from sklearn.metrics import confusion_matrix, accuracy_score, classification_report
from sklearn.linear_model import LogisticRegression
from sklearn.neighbors import KNeighborsClassifier
from sklearn.svm import SVC
from sklearn.tree import DecisionTreeClassifier
from sklearn.neural_network import MLPClassifier
from sklearn.ensemble import RandomForestClassifier
```

Figure 10. Scikit-learn for ML models and evaluation metrics

Scikit-learn offers a variety of machine learning algorithms for classification, regression, clustering, dimensionality reduction, and model selection. It is designed to interoperate with the Python numerical and scientific libraries NumPy and SciPy. The use of *sklearn.model_selection()* to split the dataset into training and test samples before training the model. The predictive model evaluates its performance using K-fold cross-validation (See Figure 11). The dataset is partitioned into k subsets or folds. The model undergoes k-fold cross-validation, where it is trained and assessed k times, with a distinct fold serving as the validation set for each iteration. The model's generalization performance is estimated by averaging the performance indicators from each fold.

```
for name, model in models:
    kfold = KFold(n_splits=num_folds, random_state=None)
    cv_results = cross_val_score(model, X_train, Y_train, cv=kfold, scoring=scoring)
    results.append(cv_results)
```

Figure 11. k-fold cross validation

The model is then trained on the dataset with accuracy_score as the performance metric for the training process as shown in Figure 12.

```
# Full Training period
 res = model.fit(X_train, Y_train)
 train_result = accuracy_score(res.predict(X_train), Y_train)
 train_results.append(train_result)
```

Figure 12. Model Training

The model prediction is discussed in Chapter 6 with the experimental results.

5.3 Azure MLHOPs Pipeline:

The following pseudocode explains the implementation of the Azure ML Pipeline.

```
config = PipelineConfiguration() #load configuration
workspace = Workspace.from_config(config) #AzureML Workspace
#Create Environment and install dependencies
env = Environment(name="train-env")
env.python.user_managed_dependencies = True
env.register(workspace=workspace)
# Create compute target
compute = ComputeTarget.attach(workspace, config.compute_name)
# Create Azure ML pipeline – a sample single step pipeline
pipeline = Pipeline(workspace=workspace, steps=[
    PythonScriptStep(
        name="data_prep",
        script_name="data_prep.py",
        inputs=[dataset.as_named_input("input_data")],
        compute_target=compute,
        environment=env)
# Submit Azure ML pipeline
run = Experiment(workspace, "AzureML-pipeline").submit(pipeline)
run.wait_for_completion()
```

This research carried out the deployment of mental health analysis using Twitter data as an MLHOPs deployment with three steps in the pipeline: *preprocess, model train and prediction.* The pipeline is created, validated, and submitted for execution. The pipeline execution is shown in Figure 13 with data inputs to the preprocessing pipeline and training pipeline.

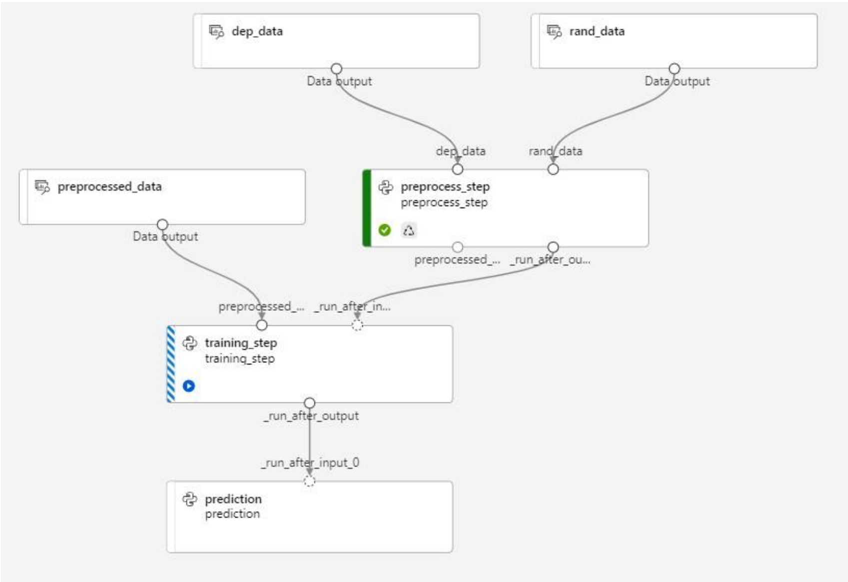

Figure 13. MLOps pipeline in execution

6 Experimental Results

Accuracy, precision, F1-score, support scores, and assessment metrics are used to convey the findings from the experiment. The ratio of accurate predictions to the total number of input samples is known as accuracy. The ratio of accurately anticipated positive observations to all expected positive observations is known as precision. Recall is the ratio of correctly predicted positive observations to all observations in actual class F1 score is the harmonic mean of precision and recall, which is a useful metric when the classes are imbalanced. The support score in a classification report is the number of actual occurrences of the class in the specified dataset. It indicates how frequently each class appears in the dataset and is employed to identify issues with the results evaluation procedure. Table-I presents the comparison of the three proposed ML models: RF, KNN and SVM performance metrics for both the prediction classes, 1 – Depressive, 0 – Non-depressive.

Table-I: Comparison of evaluation metrics for Class-1 prediction (Depressive)

Algorithm	**Accuracy**	**Precision**	**Recall**	**F1-Score**	**Support**
KNN	0.79	0.85	0.73	0.79	6300
SVM	0.85	0.82	0.86	0.84	5160
RF	0.82	0.76	0.84	0.80	4883

Table-II: Comparison of evaluation metrics for Class-0 prediction (non-depressive)

Algorithm	Accuracy	Precision	Recall	F1-Score	Support
KNN	0.79	0.73	0.85	0.79	5373
SVM	0.85	0.88	0.85	0.87	6513
RF	0.82	0.87	0.81	0.84	6790

From Table-I and II, it can be deduced that the SVM model performed well in all the evaluation metrics for both the prediction classes. RF model performed satisfactorily while KNN model performed poorly for classes. SVM model also had the best accuracy score of 85% in comparison to 79% for KNN and 82% for RF.

6.1 Model Prediction

The real-time prediction of mental health analysis model for depression is done by loading the trained model and make it determine if the test phrases are depressed or not. The test sentence can be modified to verify the authenticity and reliable prediction behaviour of the model.

Figure 14. Model Prediction

```
test_tweet = "I am low on energy levels today"
#test_tweet = "I have been awarded schlorship for my post-graduate program"

corpus = tweets_cleaner(test_tweet)

corpus

['low energy level today']

## word-embedding
test = pd.np.array([pd.np.array([token.vector for token in nlp(s)]).mean(axis=0) * pd.np.ones((300))
                              for s in corpus])

labels_pred = clf.predict(test)

result = labels_pred[0]

if result == 1:
    print("Your emotion levels show DEPRESSIVE Behavior")
else:
    print("Your emotion levels are POSITIVE")

Your emotion levels show DEPRESSIVE Behavior
```

7 Conclusion

This research tries to answer the research question in finding out the effectiveness of MLOps using social media information like Twitter in identifying the depressive behavior of the users using NLP and ML models to analyze the data and derive the conclusion on its performance. To this extent, it can be summarized that the commissioning of NLP and ML models has a considerable impact in detecting depression among the youth population, the substantial users of social media. This study has tried to achieve the reliable detection of depression and mental health status of the social media users using language models like spaCy to filter out the offensive words, and words with strong negative intonations from the Twitter dataset corpus and provide a well-processed dataset for further analysis by the ML models. The supervised

ML modes used in this research, RF, KNN and SVM performed well after the execution of the Azure pipeline where the additional dataset random_tweets.CSV was deployed and integrated, there were also found to be a few instances where they failed to correctly detect the proper emotion. This is acceptable because the dataset, though extensive and encompassing, cannot be compared with other LLMs that Google or OpenAI uses. As explained earlier this research was also designed as an attempt to make this Twitter depression prediction analysis a deployable MLOps model with end-to-end deployment with services. However, because of a few obstacles encountered during the final deployment process on Azure ML pipelines, it was only partially feasible and successful. With more advanced deep learning models and improved MLOps procedures, this study can be further enhanced.

References

Abd Rahman, R., Omar, K., Noah, S.A.M., Danuri, M.S.N.M. and Al-Garadi, M.A., 2020. Application of machine learning methods in mental health detection: a systematic review. *Ieee Access*, *8*, pp.183952-183964.

Abd Rahman, R., Zaini, F.H.M., Danuri, M.S.N.M. and Amin, A., 2022, November. The Sentiment Analysis on Mental Health Awareness by Non-Governmental Organisation's Twitter. In *2022 International Visualization, Informatics and Technology Conference (IVIT)* (pp. 185-190). IEEE.

AminiMotlagh, M., Shahhoseini, H. and Fatehi, N., 2022. A reliable sentiment analysis for classification of tweets in social networks. *Social Network Analysis and Mining*, *13*(1), p.7.

Berryman, C., Ferguson, C.J. and Negy, C., 2018. Social media use and mental health among young adults. *Psychiatric quarterly*, *89*, pp.307-314.

Chakraborty, K., Bhatia, S., Bhattacharyya, S., Platos, J., Bag, R. and Hassanien, A.E., 2020. Sentiment Analysis of COVID-19 tweets by Deep Learning Classifiers—A study to show how popularity is affecting accuracy in social media. *Applied Soft Computing*, *97*, p.106754.

Chaudhary, M., Kosyluk, K., Thomas, S. and Neal, T., 2023. On the use of aspect-based sentiment analysis of Twitter data to explore the experiences of African Americans during COVID-19. *Scientific Reports*, *13*(1), p.10694.

Chowdary, M.N., Sankeerth, B., Chowdary, C.K. and Gupta, M., 2022, August. Accelerating the Machine Learning Model Deployment using MLOps. In *Journal of Physics: Conference Series* (Vol. 2327, No. 1, p. 012027). IOP Publishing.

Conway, M. and O'Connor, D., 2016. Social media, big data, and mental health: current advances and ethical implications. *Current opinion in psychology*, *9*, pp.77-82.

Coyne, S.M., Rogers, A.A., Zurcher, J.D., Stockdale, L. and Booth, M., 2020. Does time spent using social media impact mental health?: An eight year longitudinal study. *Computers in human behavior*, *104*, p.106160.

De Choudhury, M., Gamon, M., Counts, S. and Horvitz, E., 2013. Predicting depression via social media. In *Proceedings of the international AAAI conference on web and social media* (Vol. 7, No. 1, pp. 128-137).

Finch, K.C., Snook, K.R., Duke, C.H., Fu, K.W., Tse, Z.T.H., Adhikari, A. and Fung, I.C.H., 2016. Public health implications of social media use during natural disasters, environmental disasters, and other environmental concerns. *Natural Hazards*, *83*, pp.729-760.

Gao, J., Zheng, P., Jia, Y., Chen, H., Mao, Y., Chen, S., Wang, Y., Fu, H. and Dai, J., 2020. Mental health problems and social media exposure during COVID-19 outbreak. *Plos one*, *15*(4), p.e0231924.

Garg, S., Pundir, P., Rathee, G., Gupta, P.K., Garg, S. and Ahlawat, S., 2021, December. On continuous integration/continuous delivery for automated deployment of machine learning models using mlops. In *2021 IEEE fourth international conference on artificial intelligence and knowledge engineering (AIKE)* (pp. 25-28). IEEE.

Govindasamy, K.A. and Palanichamy, N., 2021, May. Depression detection using machine learning techniques on twitter data. In *2021 5th international conference on intelligent computing and control systems (ICICCS)* (pp. 960-966). IEEE.

Kanimozhi, T. and Belina VJ Sara, S., 2022, November. Classification of Tweet on Disaster Management Using Random Forest. In *International Conference on Advancements in Smart Computing and Information Security* (pp. 180-193). Cham: Springer Nature Switzerland.

Karim, F., Oyewande, A.A., Abdalla, L.F., Ehsanullah, R.C. and Khan, S., 2020. Social media use and its connection to mental health: a systematic review. *Cureus*, *12*(6).

Kreuzberger, D., Kühl, N. and Hirschl, S., 2023. Machine learning operations (mlops): Overview, definition, and architecture. *IEEE Access*.

Kumar, P., Samanta, P., Dutta, S., Chatterjee, M. and Sarkar, D., 2022. Feature based depression detection from twitter data using machine learning techniques. *Journal of Scientific Research*, *66*(2), pp.220-228.

Lauricella, S., Small, E. and DiDomenico, V., 2022. University Is Not for the Weak: Student Communication of Mental Health on Twitter. *Journal of Digital Life and Learning*, *2*(2), pp.27-51.

Le Glaz, A., Haralambous, Y., Kim-Dufor, D.H., Lenca, P., Billot, R., Ryan, T.C., Marsh, J., Devylder, J., Walter, M., Berrouiguet, S. and Lemey, C., 2021. Machine learning and natural language processing in mental health: systematic review. *Journal of Medical Internet Research*, *23*(5), p.e15708.

Levanti, D., Monastero, R.N., Zamani, M., Eichstaedt, J.C., Giorgi, S., Schwartz, H.A. and Meliker, J.R., 2023. Depression and anxiety on Twitter during the COVID-19 stay-at-home period in 7 major US cities. *AJPM focus*, *2*(1), p.100062.

Naseem, U., Razzak, I., Khushi, M., Eklund, P.W. and Kim, J., 2021. COVIDSenti: A large-scale benchmark Twitter data set for COVID-19 sentiment analysis. *IEEE transactions on computational social systems*, *8*(4), pp.1003-1015.

Owen, D., Collados, J.C. and Espinosa-Anke, L., 2020. Towards preemptive detection of depression and anxiety in twitter. *arXiv preprint arXiv:2011.05249*.

Pavlova, A. and Berkers, P., 2020. Mental health discourse and social media: Which mechanisms of cultural power drive discourse on Twitter. *Social Science & Medicine*, *263*, p.113250.

Pérez, A., Parapar, J. and Barreiro, Á., 2022. Automatic depression score estimation with word embedding models. *Artificial Intelligence in Medicine*, *132*, p.102380.

Prakash, A., Agarwal, K., Shekhar, S., Mutreja, T. and Chakraborty, P.S., 2021, March. An ensemble learning approach for the detection of depression and mental illness over twitter data. In *2021 8th International Conference on Computing for Sustainable Global Development (INDIACom)* (pp. 565-570). IEEE.
Prieto, V.M., Matos, S., Alvarez, M., Cacheda, F. and Oliveira, J.L., 2014. Twitter: a good place to detect health conditions. *PloS one*, *9*(1), p.e86191.
Qasem, M., Thulasiram, R. and Thulasiram, P., 2015, August. Twitter sentiment classification using machine learning techniques for stock markets. In *2015 International Conference on Advances in Computing, Communications and Informatics (ICACCI)* (pp. 834-840). IEEE.
Schnebly, J. and Sengupta, S., 2019, January. Random forest twitter bot classifier. In *2019 IEEE 9th Annual Computing and Communication Workshop and Conference (CCWC)* (pp. 0506-0512). IEEE.
Tadesse, M.M., Lin, H., Xu, B. and Yang, L., 2019. Detection of suicide ideation in social media forums using deep learning. *Algorithms*, *13*(1), p.7.
Tejaswini, V., Babu, K.S. and Sahoo, B., 2022. Depression Detection from Social Media Text Analysis using Natural Language Processing Techniques and Hybrid Deep Learning Model. *ACM Transactions on Asian and Low-Resource Language Information Processing*.
Valdez, D., Ten Thij, M., Bathina, K., Rutter, L.A. and Bollen, J., 2020. Social media insights into US mental health during the COVID-19 pandemic: Longitudinal analysis of Twitter data. *Journal of medical Internet research*, *22*(12), p.e21418.
Vigo, D., Thornicroft, G. and Atun, R., 2016. Estimating the true global burden of mental illness. *The Lancet Psychiatry*, *3*(2), pp.171-178.
World Health Organization, 2020. *Mental health and psychosocial considerations during the COVID-19 outbreak, 18 March 2020* (No. WHO/2019-nCoV/MentalHealth/2020.1). World Health Organization.