Mental Health Diagnosis and Treatment Monitoring

Presenters:
Mai Tah Lee, Annie Purisch, Seth Mlodzik, Tianxing Liu

Contents

01

Introduction

Dataset Overview

Includes 17 features.

Synthetic dataset with patient demographics, symptoms, and treatment details.

Target outcome:
        Outcome categorized as **Improved**, **No Change,** or **Deteriorated**.

Numerical Features: Age, Symptom Severity, Mood Score, Sleep Quality, etc.

Categorical Features: Gender, Diagnosis, Medication, Therapy Type, etc.

0   Patient ID
1   Age
2   Gender
3   Diagnosis
4   Symptom Severity (1-10)
5   Mood Score (1-10)
6   Sleep Quality (1-10)
7   Physical Activity (hrs/week)   int64
8   Medication
9   Therapy Type
10 Treatment Start Date
11 Treatment Duration (weeks)
12 Stress Level (1-10)
13 Outcome
14 Treatment Progress (1-10)
15 AI-Detected Emotional State
16 Adherence to Treatment (%)

Our Question

How can we best predict a mental health patient's treatment outcome using their demographics, symptom severity, mood scores, and treatment types?

02
Methods

Methods - Preprocessing

Key Steps:

- Standardized numerical variables using StandardScaler.
- Encoded categorical variables using LabelEncoder and OneHotEncoder.
- Created binary features:
    - **Severity:** Symptom Severity > 5 = severe.
    - **Mood:** Mood Score ≥ 6 = good mood.
    - etc.
- Split dataset into training (80%) and testing (20%) subsets.

*Key Challenge:
Imbalanced classes (Improved: 170, No Change: 159, Deteriorated: 171).

Feature Selection

- We used Random Forest Classification to find the best combination of variables in predicting outcome
- The best combination was:
  - `'Age', 'Symptom Severity (1-10)', 'Mood Score (1-10)','Physical Activity (hrs/week)', 'Treatment Duration (weeks)''Treatment Progress (1-10)'`
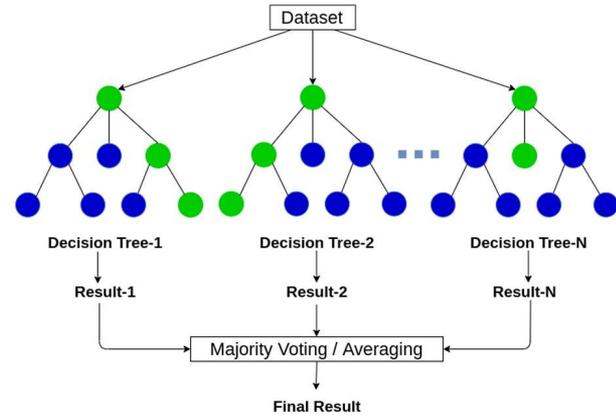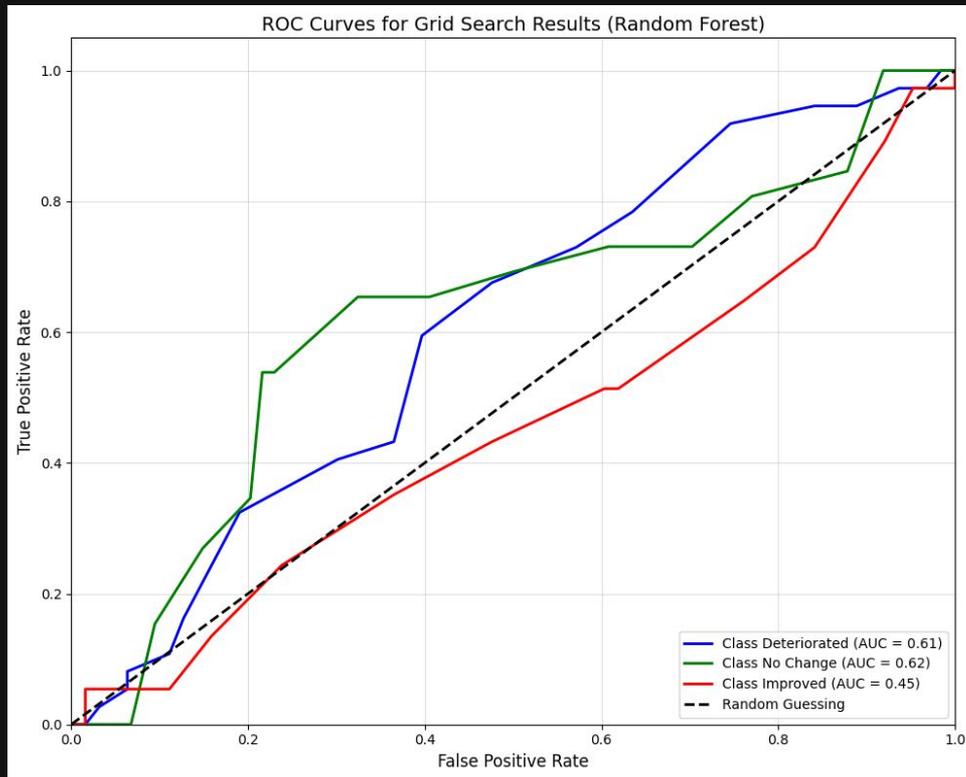
Models Applied

- Best Cross-Validated Accuracy: **0.32**

- Logistic Regression Accuracy: **0.37**

- Decision Tree Accuracy: **0.32**

- Random Forest Accuracy: **0.50**

Model Comparison: Accuracies
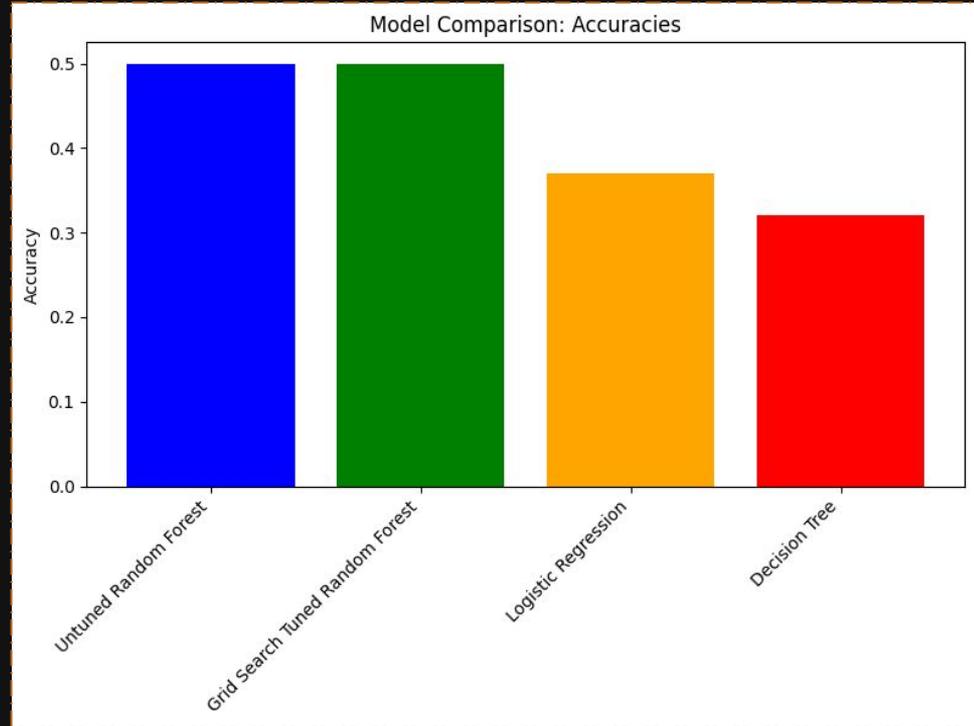
Random Forest

Prediction Accuracy of Outcome

- 'Deteriorated': AUC=**.61**

- 'No Change': AUC=**.62**

- 'Improved': AUC=**.45**

ROC Curves for Grid Search Results (Random Forest)

— Class Deteriorated (AUC = 0.61)
— Class No Change (AUC = 0.62)
— Class Improved (AUC = 0.45)
- - - Random Guessing

03

Outcome Highlights

Performance Metrics

- We split the dataset into training (80%) and testing (20%) subsets to evaluate the performance of three models: Logistic Regression, Decision Tree, and Random Forest.

- Each model was trained on the training data and tested on the unseen test set.

- Key performance metrics—Accuracy, Precision, Recall, and F1-Score—were calculated to assess prediction quality, resulting in values like Accuracy
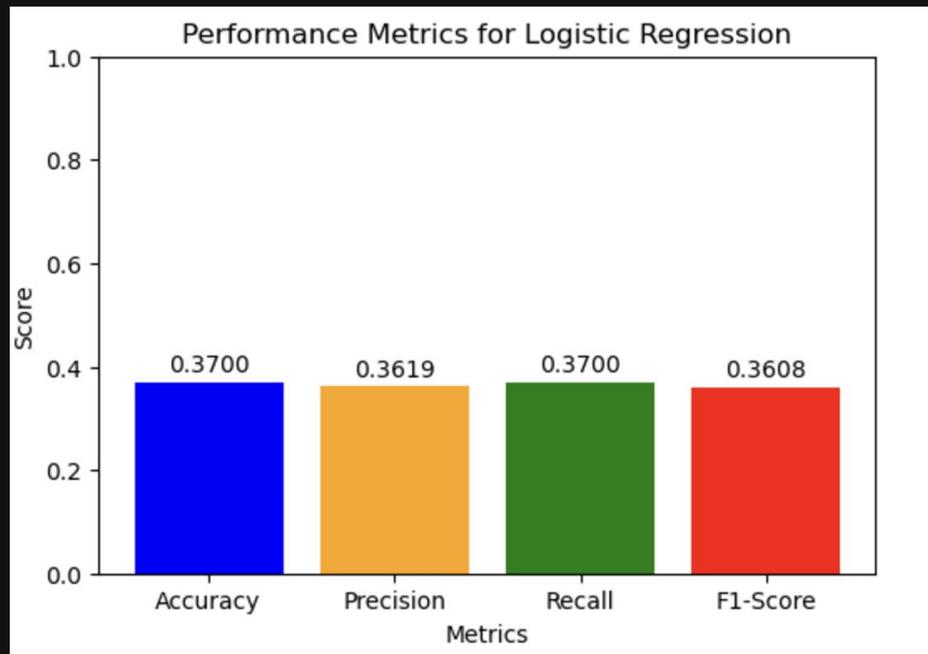
Logistic Regression Metrics

Accuracy: 0.3700
Precision: 0.3619
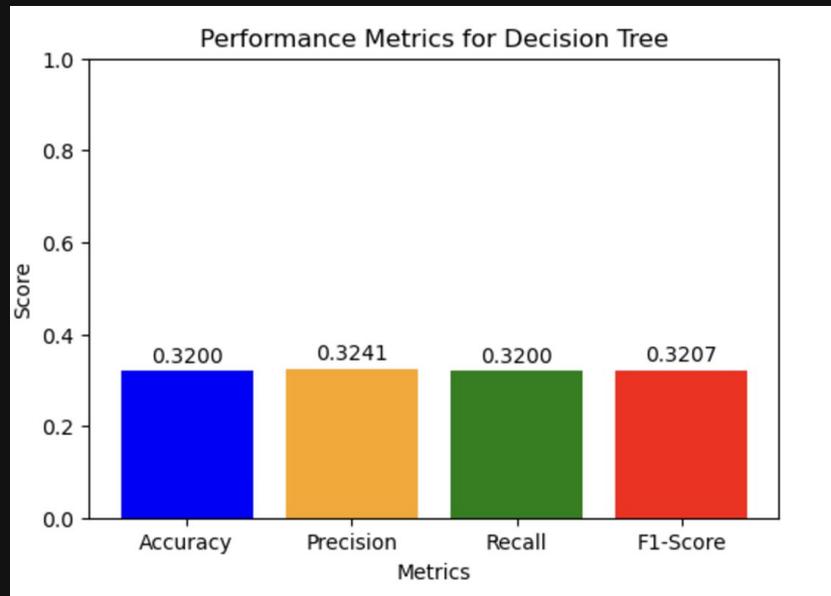Recall: 0.3700
F1-Score: 0.3608

Decision Tree Metrics

Accuracy: 0.3200
Precision: 0.3241
Recall: 0.3200
F1-Score: 0.3207

Performance Metrics for Decision Tree
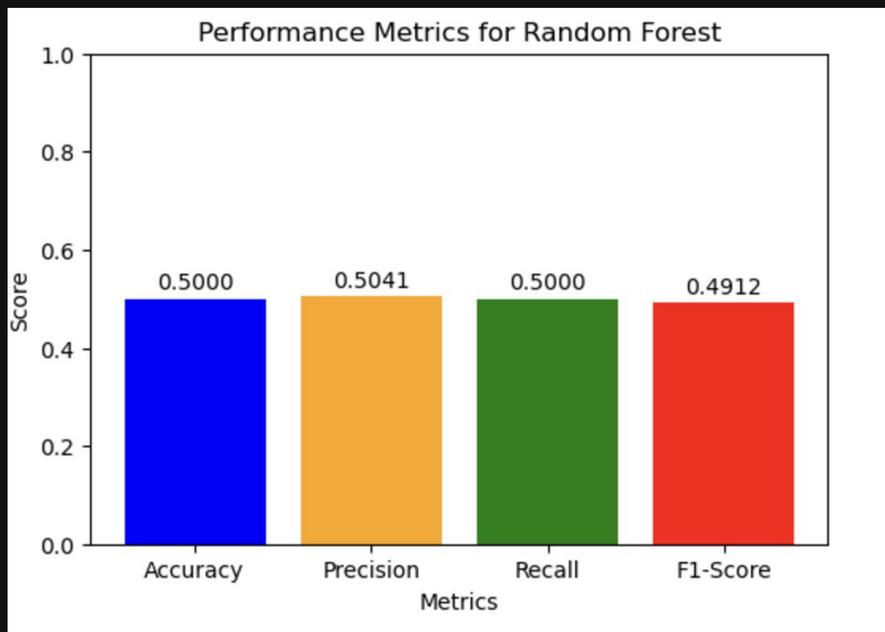
Random Forest Metrics

Accuracy: 0.5000
Precision: 0.5041
Recall: 0.5000
F1-Score: 0.4912

- Conclusion: Random Forest achieved the highest performance (Accuracy: 0.5000, F1-Score: 0.4912), followed by Logistic Regression and Decision Tree.

Performance Metrics for Random Forest

04

Limitations/Challenges

Limitations & Challenges

- Challenges in evaluating logistic regression due to non-binary outcome categories ("Deteriorated," "No Change," and "Improved").

- Difficulty converting outcomes to binary variables, leading to imbalanced datasets when ignoring certain outcomes.

- Imbalance in predictions when focusing on subsets of outcomes (e.g., excluding "Deteriorated" or "Improved").

- Other things to consider is that the data is synthetic consisting of only 500 rows. There may be some skewed data.

Thank You!

Q&A

Sources

https://www.kaggle.com/datasets/uom190346a/mental-health-diagnosis-and-treatment-monitoring