

# MyGraine: Predicting Migraines Through Various Machine Learning Models Utilizing User-Inputted Data

Rebecca S. Zhu, Rucha Dave

Nashua High School South, 36 Riverside Street, Nashua, NH, 03062, U.S.A.; rebecca3@mit.edu

**ABSTRACT:** Migraines are the sixth most disabling illness in the world. Although they are often seen as just headaches, they severely lower the quality of life for afflicted patients. Currently, migraine monitoring applications are ineffective and inaccurate. Through this project, multiple types of machine learning models were trained to predict the occurrence of migraines in individuals. Data that was accumulated from an online public database was converted into a usable format. In order to reduce imbalance and bias in the database, we used Synthetic Minority Over-Sampling Technique (SMOTE), a data augmentation technique on the database. Data visualization enabled patterns to be observed, determining the most significant factors for a migraine occurrence. Two models were trained to accurately predict the occurrences of migraines: a logistic regression and a random forest. The logistic regression model achieved significant results with a 97% accuracy, but the random forest was also close with an 89% accuracy. The number of triggers was found to be the deciding factor in determining migraine occurrences. The second most important factor was the number of helping factors present, such as getting enough sleep or exercise. Finally, a user-friendly website was developed. This website took in biometrics and user-inputted data to display the probability of getting a migraine, using the logistic regression model.

**KEYWORDS:** Migraines; Predicting Migraines; Migraine Triggers; Machine Learning; SMOTE; Feature Selection; Logistic Regression Model; Random Forest Model.

## ■ Introduction

Migraines are the sixth most disabling and third most prevalent illness in the world. More than four million people suffer from chronic migraines at least fifteen times a month.<sup>1</sup> A migraine is a recurrent headache that causes moderate to severe pain experienced through throbbing or pulsing in the head. Often, it is also associated with sensitivity to light and loud sounds, nausea, and vomiting.<sup>2</sup> Those who suffer from migraines are unable to complete daily tasks. But while most research is focused on finding treatments and cures for migraines, there is a need to help patients predict their future migraine occurrences from their migraine triggers, which would allow them to avoid triggers and minimize pain.

Some companies have tried to create “migraine diaries” that allow users to track their migraines by inputting migraine occurrences. However, there are two main problems with these efforts: they are often time-consuming for patients, and they do not allow users to see the causes of their migraines, meaning that they cannot take preventative measures. Migraine Buddy© and Migraine Coach©, for example, are two popular migraine tracking apps that allow patients to record how long their migraines last and how frequently they occur. Although Migraine Buddy allows patients to input factors that may have caused their migraines, the process is intensive and takes a lot of time to complete every day. Some information that Migraine Buddy tracks is “the type of attack, pain level, your location when your migraine began, and potential triggers”.<sup>3</sup> Though they attempt to improve patient lifestyle, these apps

are not designed to predict the likelihood of a future migraine occurrence based on the current status of a user; they are only meant to be used as a calendar to store information on past migraines. Additionally, this only allows for retrospective information to be relayed to the individuals and no information about preventing future migraines.

To address this gap, we built machine learning models that predict migraines based on individual biometric and user-inputted data. Machine learning methods can differentiate between a migraine occurrence and a non-migraine occurrence after being trained with data. Machine learning methods can be applied to find correlations between certain migraine triggers and the probability of a migraine occurrence. Our machine learning models take user-inputted data to predict migraine occurrences. Users also have the option to enter personal triggers through the website we built. Once a user confirms or inputs a new migraine, the model categorizes migraines based on the inputted triggers. In the future, patients will have the option to provide their data to physicians, who will be able to use it to improve treatment.

Predicting migraine occurrences is crucial in both its prevention and mitigation. Patients can be more aware of their triggers and avoid them, sparing them the pain of severe migraines. Doctors and other medical providers can analyze the data to change and administer treatments per their discretion. Our solution specifically tackles the problem of the lack of an accurate prediction system for migraines. By implementing a machine learning model that can find patterns between indi-

vidual migraine factors inputted into the application, we hope to be able to improve the lifestyle of countless patients who suffer from these severe headaches.

In a prior study,<sup>4</sup> the authors created their own dataset from over half a million migraines. They reported the most common triggers, but they did not fit predictive models to do so. In another study, researchers created an application that provided video, audio, and animated instructions about learning how to manage migraines.<sup>5</sup> Additionally, they provided a headache diary that tracked triggers that users self-report. But because these were self-reported by patients, they were susceptible to biases. Users may not have had enough knowledge on how migraines occur and might have mistaken coincidences for trigger factors. Because migraines involve physical pain, patients may not have the awareness to remember and record details about their migraines. By incorporating machine learning into the process, human errors become less common and the results more reliable. In another research study,<sup>6</sup> the authors administered a questionnaire and used the results to train a Naive Bayes classifier and Support Vector Machine (SVM). Additionally, they reported that feature selection using the Naive Bayes improved the prediction accuracy from 67% to 93%, while feature selection improved the accuracy from 90% to 95% for the SVM. Other researchers have used deep learning models and wearable devices to predict migraines.<sup>7</sup> For many migraine studies using machine learning, researchers analyzed electroencephalograms (EEG) patterns in their data. One study utilized an ANOVA test, a hypothesis test determining if something is different and statistically significant, to classify EEG patterns.<sup>8</sup> They also utilized a neural network and achieved a total accuracy of 90.9%. In another study, researchers used AR Burg, which was an autoregression based method, and MUSIC method, which was a subspace-based model, and SVM to analyze migraine EEG signals.<sup>9</sup>

In this study, we addressed the open problem of how machine learning models can be utilized to predict migraines. Our study is different from previous ones because we focused on mitigating the imbalance in the data before applying machine learning techniques in order to improve the accuracy of the machine learning models. We also used different machine learning models, specifically a Logistic Regression model and a Random Forest model, to obtain our results. Instead of focusing on EEG activity, which requires a professional, we used common migraine triggers that many individuals have to predict future migraine occurrences, making our implemented model accessible to patients from their homes.

## ■ Results and Discussion

### Data Collection

The database we used was called “Analysis of Trigger Factors in Episodic Migraineurs Using a Smartphone Headache Diary Applications”.<sup>10</sup> This database was published on February 22, 2016. The study that published this database was done by Jeong-Wook Park, Min Kyung Chu, and other collaborating scientists and published on Plos One. In the dataset, there were numerous migraine triggers or features that could increase a chance of a migraine, such as “stress”, “excess sleep”, “exercise”,

“fatigue”, “drinking”, “overeating”, “caffeine”, “smoking”, “travel”, etc. These features were used to predict ‘getting a migraine’ or ‘not getting a migraine’, which were the two classes for classification. Each feature had two options for selection and the answers were represented in a binary. Responding “yes” to a feature equated to 1, while “no” equated to 0. For example, if a user had stress, then the database had a 1 under the stress feature. Each feature represented a column in the database while each user and his or her migraine prediction was a row in the database. The dataset contained 4580 entries including both migraine and non-migraine days.

### Proposed Mechanism

Recently, machine learning has been used to classify and predict many diseases and illnesses such as COVID-19. Due to machine learning’s ability to learn the representation of the data through the models, it was often used to classify labels for given data. Because our question focused on building a model that classified large amounts of data into categories accurately, we decided that machine learning would improve the prediction accuracy for migraines in our research. Machine learning enabled computers to find patterns through numerous data samples in order to predict the outcome of new data. The two types of machine learning models were classification and regression. Classification allows data to be put into discrete outcomes. Regression, in contrast, is a continuous model that can predict probabilities of a continuous target. The two models used in our research were a logistic regression model and a random forest model. Although the logistic regression model was most often used as a regression model, we repurposed it for classification to match our needs. The random forest model, which can be used for both purposes, performed as a classification model in our research as well. Figure 1 shows the process used from getting the data to training the model. It represents the process for both the logistic regression and the random forest model.

It represents the process for both the logistic regression and the random forest model.

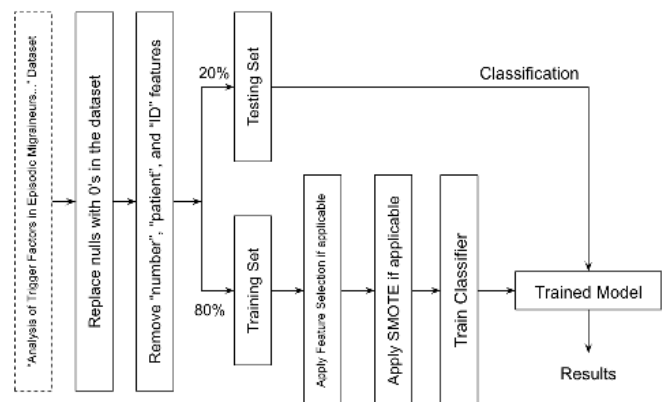


Figure 1. The proposed approach for our machine learning models to classify migraines.

To have a basis of comparison between both models, we had certain parameters that were defined as constants. We reserved 30% of the dataset for our testing data, while the other

70% became the training Synthetic Minority Over-Sampling Technique (SMOTE)<sup>11</sup> was only implemented it on the training data. We used the testing data to get our results in accuracy, precision, and macro-average F1 scores. Macro-average scores were calculated so that they did not take class imbalance into consideration. Thus, we directly compared our results from the machine learning models with and without SMOTE. SMOTE is a data augmentation method used to increase the number of under sampled instances of a class using the “k-nearest neighbors” algorithm. In our case, the under sampled instances were “no migraine” occurrences. Because of the significant imbalance observed in the dataset, we utilized the imbalanced-learn library in Python to applying SMOTE to our dataset.<sup>12</sup>

Table 1. Metrics for the discussion of the results.

Abbreviations	Descriptions
TP: True Positive	The number of patients predicted by the machine learning models as migraineurs that were migraineurs.
TN: True Negative	The number of patients the machine learning models predicted will not get a migraine occurrence that did not get a migraine occurrence.
FP: False Positive	The number of patients predicted by the machine learning models as migraineurs but did not actually suffer a migraine.
FN: False Negative	The number of patients that the machine learning models predicted will not get a migraine but got a migraine.

In evaluating our results, we did not only focus on accuracy, as accuracy alone could not provide detailed insight about our conclusions. Instead, we evaluated macro-average precision, macro-average recall, and macro-average F1 scores. In Table 1, we describe the metrics that were used to evaluate the results of our research. These metrics were also used in the formulae for accuracy, macro-average precision, macro-average recall, and macro-average F1 scores, which are shown below in Equation 1:

$$\text{Equation 1: } \text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}, \text{Precision} = \frac{TP}{TP + FP}, \text{Recall} = \frac{TP}{TP + FN}, F1 = \frac{2 * \text{precision} * \text{recall}}{\text{precision} + \text{recall}}$$

In our experiments, we used the accuracy, precision, recall, and macro-average F1 scores from the Scikit-Learn open-source Python library.<sup>13</sup> Additionally, these metrics were evaluated with the testing data to determine whether our models performed well on data they had never seen before. In other words, the higher the metrics were, the better our models generalized to new data.

### Data Visualization

Before training our machine learning models, we visualized our data using different Python libraries to gain an understanding of overall patterns and trends in the data. In Figure 2, the class imbalance in the dataset was evident; there were over 3 times as many “no migraine” occurrences as there were “migraine” occurrences. In Figure 3, the correlation between the number of




Figure 2. The comparison of migraine occurrences and the bias in the dataset

triggers an individual suffered and whether they got a migraine is observed. We saw that there was an association between the number of triggers and the migraine occurrence: as one got more triggers, it became more probable that one will also get a migraine. There are a few discrepancies in Figure 3. The proportion of no migraine occurrences goes to 0 for 7 triggers and then rises back up to 0.5 for 8 and 9 triggers. While this may seem like a contradiction to the claim above, it is explained by the dataset we used. Our database only had 2-3 values for each of those categories (7, 8, and 9 triggers). Because of this limited size of data, those sections of the graph are not reliable. Our machine learning models were able to take advantage of these associations in order to accurately predict when a user will get a migraine.

### Logistic Regression

One machine learning model utilized in this research was a logistic regression. The logistic regression model is a statistical model that calculates the probability of two different results given one or more predictors or features. The outcome of our logistic regression was a dichotomous variable, meaning that there were only two possible outcomes. In predicting migraines, a migraine occurrence corresponded to 1 while a no migraine occurrence corresponded to 0. Given certain triggers of migraines, the logistic regression algorithm found a pattern between the triggers and predicted whether a migraine will occur or not. The probability of the migraine occurrence or  $P(x)$  for an observation with feature  $x$  was determined by the logistic equation, otherwise known as a sigmoid function, which was denoted by Equation 2:

$$\text{Equation 2: } P(x) = \frac{1}{1 + e^{-(\beta_0 + \beta_1 x)}}$$

where  $\beta_0$  and  $\beta_1$  represented parameters to the sigmoid function as determined by the individual features and data of the migraine dataset.  $X$  represented a data matrix given by the inputs and features of our dataset.

While the logistic function provided a probability that represented a continuous variable, we used the probability to determine whether an individual experienced a migraine occurrence or not. When training on the dataset without SMOTE




Figure 3. The stacked bar chart of the number of triggers versus migraine occurrence.




Figure 4. The confusion matrix for the logistic regression model trained on the testing data without SMOTE.

but with all the features, the logistic regression achieved an accuracy of 0.95 and a macro-average F1 score of 0.88. The confusion matrix for the logistic regression model is shown in Figure 4. We saw that generally the model performed well predicting real no migraines while it inaccurately predicted no migraines when the individual did experience a migraine. To quantify, it accurately predicted 1193 instances of no migraines as no migraines and 119 instances of migraines as migraines. However, it inaccurately predicted 46 actual migraine occurrences as no migraines and 16 actual no migraine occurrences as migraine occurrences.




Figure 5. The confusion matrix for the logistic regression model trained on the testing data with SMOTE, an improvement from the classification without SMOTE.

In order to see if the imbalance in the classes caused the false negatives, we implemented SMOTE, a method of data augmentation, on the data. SMOTE produced synthetic instances of the migraine class because there were more instances of the no migraine class than the migraine class. Without SMOTE, it seemed obvious that the model could have kept predicting no migraine and kept getting a high accuracy just because of

the high ratio of the no migraine occurrences to the migraine occurrences. To see if the model could improve its classification on migraine occurrences, SMOTE must be utilized to mitigate the imbalance in the dataset. After applying SMOTE, there were 2832 instances of the migraine class and no migraine class each. The model with SMOTE outperformed the model without SMOTE, as its accuracy improved by 2% to 97%, and its macro-average F1 score improved by 0.06 to 0.94. The confusion matrix in Figure 5 detailed the TP, TN, FP, and FN classifications that the model achieved on the testing data. We saw that the number of FN, or when the model predicted no migraine, but the individual did have a migraine, fell to 0 as opposed to the 46 from the model without SMOTE. However, after utilizing SMOTE, we saw that the false positive rate increased. To quantify, the model predicted 16 no migraines as migraines without SMOTE, but predicted 39 no migraines as migraines with SMOTE. The increase in FP and the decrease in FN is due to the newly formed balanced dataset from SMOTE using the same logistic regression model as before. Although the false positive rate increased, the macro-average F1 score still increased and the accuracy increased with SMOTE because of the large improvement in lowering the number of false negatives. Thus, Figure 5 shows that SMOTE enhanced our model's accuracy and macro-average F1 score. Additionally, SMOTE suggests that the model got better at predicting true migraine occurrences.

The results for the model with and without SMOTE are shown below in Table 2. We saw that the macro-average precision score decreased due to the increase in FP. However, the macro-average recall increased more than the decrease in macro-average precision, which resulted in an overall increase in the macro-average F1 score. The increase in the number of TP, which was when the model predicted correctly for a migraine occurrence, ultimately increased the overall accuracy for the logistic regression with SMOTE.

Table 2. The results obtained with and without SMOTE on the testing data with the logistic regression model.

Colony Diameter (mm)	Sound Wave Frequency (Hz)			
	0	432	802	4000
T1.1	2.198	2.140	2.189	1.634
T1.2	2.185	1.736	1.754	1.993
T1.3	2.273	2.315	1.536	1.185
T2.1	2.650	1.677	1.988	1.237
T2.2	2.076	2.144	1.610	1.485
T2.3	2.031	1.795	1.969	1.384
T3.1	3.733	2.291	1.590	1.447
T3.2	5.395	1.676	1.490	1.689
T3.3	2.055	2.109	1.137	0.920
Average	2.733	1.987	1.696	1.442

The receiver operating characteristic (ROC) curve shown in Figure 6 details the opportunity cost between the false positive and true positive rates. The ROC curve was used to determine the logistic regression's performance on distinguishing between what truly was a migraine occurrence and what was not. The baseline curve, represented by the dashed red line, shows the performance of an untrained and inaccurate classifier, as the true positive rate was equivalent to the false



positive rate. When the curve was closer to the upper left corner, the classifier became more accurate. This curve represented the performance of the logistic regression with SMOTE on the dataset. As shown by the solid blue line, which represents the logistic regression's ROC curve, the logistic regression was very accurate at predicting migraines. This means the logistic regression classifier was able to correctly predict migraine occurrences almost all the time. Overall, from the ROC curve, our model performed well on the testing data.

### Random Forest

The second model that was created in this research was a random forest. The random forest model is a supervised regression-based machine learning algorithm that predicts outcomes based on a step-by-step decision-making process. A decision tree is an intuitive branching system that leads each entry to a conclusion based on similarities and differences.




Figure 6. The receiver operating characteristic (ROC) curve for the logistic regression model with SMOTE applied on the testing data

In the random forest model, many uncorrelated decision trees predicted outcomes individually. These results were compared through ensemble learning and the random forest model reached a conclusion. This prevented any individual errors or outliers from influencing the final answer and safeguarded against overfitting, which is one of the larger errors of machine learning. Overfitting is the phenomena when the model performed well on the training data but poorly on the testing data. In other words, the model did not generalize well to newly seen data. To avoid overfitting, random forest utilizes ensemble learning to reach a conclusion. In the context of this research, the model took in the various triggers and symptoms and predicted whether a migraine will occur or not. In addition, the random forest model created many different decision trees, and then viewed the overall trend from all those trees to reach a conclusion. Because of this, even if a few decision trees did overfit, the general trend would ensure that overfitting is avoided.

Before the model was tested, a baseline estimation was done. This gave us a target to reach through our model. If the model randomly chose migraine for 50% of the data and no migraine for the remaining half, we received an accuracy of 75.6%.

For the random forest to be a proper predictor of migraines, it would need to reach an accuracy significantly higher than 75.6%. This performance indicator also revealed the imbalance of the dataset as the baseline accuracy was above 50%.

Figure 7 highlights an example of a certain node in one of the decision trees from our random forest model. The top stem includes the deciding factor for all the values that fell under it. Based on Figure 7, this variable was the number of helping factors being less than or equal to 0.5. If there were more than 0.5 helping factors present, then the process moved to the right. Otherwise, it continued to the left.

The line underneath calculated the mean square error (MSE) of this node. This error was calculated by finding the average of the sum of all errors squared. Calculating the mean square error of data provided a metric to create a decision tree. For optimal results, the model minimized the MSE of all nodes. The next line gave the total number of samples that were pass-through this node, and the last line (value) gave a prediction for all values under this node. In this figure, for the first node, the prediction for all 2201 samples that pass through it was no migraine. In this way, each data value passed through the entire decision tree and reached the end, at which no further distinctions could have been made between values. As seen in Figure 7, the very bottom nodes, called the leaf nodes, provided an individualized prediction for each data value. Looking at the node at the bottom left corner, there were 228 values that fall under this specific category and the prediction for these was 0.2, or no migraine.

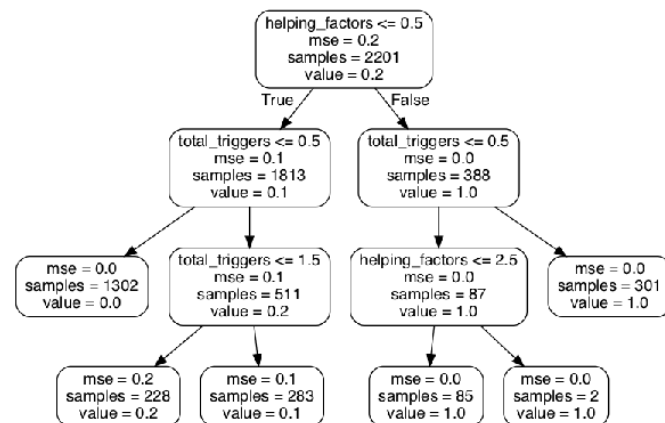


Figure 7. One section of a single decision tree from the random forest model.

Many similar decision trees were constructed in the random forest model, each leading the data to a conclusion of migraine or no migraine. The majority response was then returned. Using the testing data, our random forest model achieved a prediction accuracy of 89.66%. This was the result with SMOTE. The average absolute error was 0.1. This was a significantly higher accuracy than the baseline prediction's accuracy of 75.6% and meant that the model improved prediction accuracy. We saw that SMOTE also improves the accuracy of the classification by mitigating the imbalance in the dataset. The random forest algorithm provided a feature importance attribute that weighed certain features based on the contribution to the classification. From evaluating our

model, the most contributing feature to the classification was the number of triggers while the second most contributing feature was the number of helping factors such as the amount of exercise and sleep. While the model showed that an increase in the number of triggers contributed the most to migraines, the opposite was true for the number of helping factors. The more the helping factors, the lower the chance for a migraine. Through future research, we hope to further investigate these factors along with others to further understand these weights and their effects on migraines.

### Website Development

Research was conducted in making an application for the users to predict migraine occurrences. Because the logistic regression with SMOTE outperformed all the other models utilized in this research, we implemented the logistic regression model with SMOTE on our website. In Figure 8, the layout of the website is shown. Users were prompted with a quiz where they entered certain personal triggers through a simple questionnaire having the same features and inputs as the features used to train the logistic regression model with SMOTE. Then, in the backend of the website, we used Flask, which is a micro web framework written in Python, to load our logistic regression model, predict the probability a user has a migraine, and display the probability on the screen.



Figure 8. The layout of our MyGraine website that allows users to predict their migraines using the logistic regression model with SMOTE.

### Conclusion

Recently, there has been a demand for accurate migraine prediction through machine learning to help physicians advise patients what is best to avoid future migraines. By analyzing the dataset and utilizing SMOTE data augmentation, we provided an accurate, fast, and user-friendly application that might help migraineurs predict their future migraines. Through research and experiments, it was understood that many factors affect migraines. Both the logistic regression and random forest models achieved high results, but the logistic regression outperformed existing models using a standard dataset. To quantify, the logistic regression achieved a 97% prediction accuracy when SMOTE was employed. Through the random forest model and feature selection, it was found that the deciding factor for migraines was the number of triggers present. A website, MyGraine, was created where both the

triggers and helping factors can be entered to receive a prediction. We made the following contributions through our paper:

- We concluded that the Logistic Regression model outperformed many of the previous models for migraine classification. To quantify our success, the Logistic Regression model achieved a 97% prediction accuracy on the test data.
- We determined that feature selection could improve a Random Forest model's accuracy to 89%. By choosing different features to train our Random Forest with, we used the more discriminative features in order to improve a model's classification accuracy.
- We demonstrated that SMOTE data augmentation was useful in mitigating bias in the dataset and improved the classifiers' accuracies on an imbalanced dataset.

In the future, we want to research into federated learning<sup>14</sup> so that individuals can have their own machine learning model tailored to their own migraine triggers. Additionally, we are in contact with the American Registry for Migraine Research in order to get a larger database to confirm our results and have more extensive testing. We also want to expand our website into a mobile application through Xcode to allow patients to record migraines and the predictions to send to healthcare providers as well as advise patients about their migraine predictions. The application could be connected to Apple HealthKit, allowing the app to use biometrics as inputs as well as user-inputted features. Eventually, through testing with migraine patients with various backgrounds, our logistic regression model will become more accurate and will be able to impact the lives of many from around the world.

### Methods

A critical part of our research was in the code that we wrote. The code was essential in order to produce similar results to our findings. Our GitHub repository with all the code and resources will be available upon publication.

### Data Manipulation

In order to use the database, we pre-processed the data. There were some empty and "null" spots in the database. Here, we put 0's in these places because it was reasonable to assume that the answer was no for the question. The questions being asked for data collection were about important measures taken in a day such as exercising or sleeping in. We expected and assumed that the subjects would remember any important activities they have taken in their day. If they did not report an answer for a question, then we could consider that activity "not significant" for the day and thus assumed that either the subject did not partake in it or was not exposed to it a significant duration of time (which would mean that it most likely did not affect the migraine occurrence). Next, we dropped features such as "number", "patient", and "ID" because these were unique for every individual and were used to identify the individual. The number specifies what number entry the individual was. For example, the first person in the database, had a 1 in the "number" category, while the second person in the database had a 2 in the "number" category and so on. The "patient" feature was a unique combination of letters and numbers

given to the patient that had no correlation to whether a user experiences a migraine occurrence or not. Similarly, the "ID" feature was a random combination of numbers and letters given to each patient. Since "number", "patient", and "ID" had no correlation to whether a user experienced a migraine or not, we removed them before training our machine learning models.

### Machine Learning Models

The machine learning models utilized were implemented in Python and different Python libraries. The data was in a comma-separated values (CSV) file. In order to read the data in the file, we utilized the Pandas library to load the data. For data visualization, we employed the Seaborn and Matplotlib.pyplot library. Additionally, we used the Sklearn library for SMOTE, the logistic regression model, the random forest model, splitting the training and testing data, the confusion matrices, the ROC curves, and the decision tree visualization.

### Website

The website was coded in HTML and styled using CSS. Additionally, the backend was in Flask library in Python. In order to load the logistic regression model into the website, we first used the Pickle library in Python to save the logistic regression model. The Flask backend dealt with processing the user-inputted features in the questionnaire, predicting the probability of the migraine after the user submitted the questionnaire, and displaying the result to the user on the website.

### Acknowledgements

The authors would like to thank our teacher Mrs. Cynthia Pitkin from Nashua High School South and Professor Marek Petrik from the University of New Hampshire for their valuable feedback and help. We declare no conflict of interest regarding the publication of this research.

### References

1. Migraine Research Foundation. (n.d.). Migraine Facts. Retrieved from <https://migraineresearchfoundation.org/about-migraine/migraine-facts/> (accessed March 16, 2020)
2. Mayo Clinic Staff. (2019, November 1). Migraine - Symptoms and causes. Retrieved from <http://www.mayoclinic.org/diseases-conditions/migraine-headache/symptoms-causes/syc-20360201> (accessed March 20, 2020)
3. Migraine Buddy. (n.d.). Migraine Buddy. Retrieved from <https://migrainebuddy.com/> (accessed April 1, 2020)
4. J. Berengueres and F. Cadiou, "Migraine factors as reported by smartphone users," 2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), Orlando, FL, 2016, pp. 271-274, doi: 10.1109/EMBC.2016.7590692.
5. C. Liu, K. A. Holroyd, Q. Zhu, K. Shen and W. Zhou, "Design and implementation of a behavioral migraine management iPhone app for adolescents with migraine," 2010 IEEE International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM), Montreal, QC, 2010, pp. 1-6, doi: 10.1109/WOWMOM.2010.5534985.
6. Garcia-Chimeno Y, Garcia-Zapirain B, Gomez-Beldarrain M, Fernandez-Ruanova B, Garcia-Monco JC. Automatic migraine classification via feature selection committee and machine learning techniques over imaging and questionnaire data. BMC Med Inform Decis Mak. 2017;17(1):38. Published 2017 Apr 13. doi:10.1186/s12911-017-0434-4
7. S. Mohan and A. Mukherjee, "MigraineCloud," SoutheastCon 2018,

- St. Petersburg, FL, 2018, pp. 1-7, doi: 10.1109/SECON.2018.8478869.
8. E. Sayyari, M. Farzi, R. R. Estakhrooiech, F. Samiee and M. B. Shamsollahi, "Migraine analysis through EEG signals with classification approach," 2012 11th International Conference on Information Science, Signal Processing and their Applications (ISSPA), Montreal, QC, 2012, pp. 859-863, doi: 10.1109/ISSPA.2012.6310674.
9. S. B. Akben, A. Subaşı and M. K. Kıymık, "Comparison of sub-space based MUSIC and AR BURG methods in diagnosis of migraine by support vector machines," 2010 IEEE 18th Signal Processing and Communications Applications Conference, Diyarbakir, 2010, pp. 192-195, doi: 10.1109/SIU.2010.5651458.
10. Min Kyung Chu, J.-M. K. (n.d.). Analysis of Trigger Factors in Episodic Migraineurs Using a Smartphone Headache Diary Applications. Retrieved from <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0149577#abstract0> (accessed March 20, 2020)
11. Nitesh V. Chawla, Kevin W. Bowyer, Lawrence O. Hall, and W. Philip Kegelmeyer. 2002. SMOTE: synthetic minority over-sampling technique. J. Artif. Int. Res. 16, 1 (January 2002), 321-357.
12. Guillaume Lemaitre, Fernando Nogueira, and Christos K. Aridas. 2017. Imbalanced-learn: a python toolbox to tackle the curse of imbalanced datasets in machine learning. J. Mach. Learn. Res. 18, 1 (January 2017), 559-563.
13. Scikit-learn: Machine Learning in Python, Pedregosa et al., JMLR 12, pp. 2825-2830, 2011.
14. Google Research. (2017, April 6). Federated Learning: Collaborative Machine Learning Without Centralized Training Data. Retrieved from <https://ai.googleblog.com/2017/04/federated-learning-collaborative.html> (accessed May 15, 2020)

### Authors

Rebecca Zhu is a rising high school senior at Nashua High School South. She is interested in majoring in Computer Science on an artificial intelligence track.

Rucha Dave is also an upcoming senior at Nashua High School South. She is hoping to major in Bioengineering with a minor in Computer Science.