

A Machine Learning Analysis of COVID-19 Mental Health Data

Mostafa Rezapour (✉ rezapom@wfu.edu)

Wake Forest University

Lucas Hansen

Wake Forest University

Research Article

Keywords: COVID-19 Pandemic, Healthcare Workers, Mental Health, Supervised Learning, Unsupervised Learning, Feature Selection

Posted Date: January 7th, 2022

DOI: <https://doi.org/10.21203/rs.3.rs-1129807/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

A Machine Learning Analysis of COVID-19 Mental Health Data

Mostafa Rezapour^a, Lucas Hansen^a

^a*Department of Mathematics and Statistics, Wake Forest University, NC, U.S.*

Abstract

In late December 2019, the novel coronavirus (Sars-Cov-2) and the resulting disease COVID-19 were first identified in Wuhan China. The disease slipped through containment measures, with the first known case in the United States being identified on January 20th, 2020. In this paper, we utilize survey data from the Inter-university Consortium for Political and Social Research and apply several statistical and machine learning models and techniques such as Decision Trees, Multinomial Logistic Regression, Naive Bayes, k-Nearest Neighbors, Support Vector Machines, Neural Networks, Random Forests, Gradient Tree Boosting, XGBoost, CatBoost, LightGBM, Synthetic Minority Oversampling, and Chi-Squared Test to analyze the impacts the COVID-19 pandemic has had on the mental health of frontline workers in the United States. Through the interpretation of the many models applied to the mental health survey data, we have concluded that the most important factor in predicting the mental health decline of a frontline worker is the healthcare role the individual is in (Nurse, Emergency Room Staff, Surgeon, etc.), followed by the amount of sleep the individual has had in the last week, the amount of COVID-19 related news an individual has consumed on average in a day, the age of the worker, and the usage of alcohol and cannabis.

Keywords: COVID-19 Pandemic, Healthcare Workers, Mental Health, Supervised Learning, Unsupervised Learning, Feature Selection

1. Introduction

In late December 2019, the novel coronavirus (Sars-Cov-2) and the resulting disease COVID-19 were first identified in Wuhan China [1]. The disease slipped through containment measures [2], with the first known case in the United States being identified on January 20th, 2020 [3]. As

Email addresses: rezapom@wfu.edu (Mostafa Rezapour), HansLG18@wfu.edu (Lucas Hansen)

countries around the world grappled with the new disease, measures once considered extreme such as complete shutdowns of economies, commonplace social distancing, and near universal mask wearing have now become the norm. New data and its analysis have heavily influenced policymakers around the globe and vastly transformed our knowledge of the disease and its effects [4]. In this paper, we examine the indirect effects of COVID-19 on the mental health of frontline workers in the United States and offer approaches to help our frontline workers retain their mental health during these trying times.

The United States Center for Disease Control and Prevention (CDC) defines an epidemic as *an increase, often sudden, in the number of cases of a disease above what is normally expected in that population in that area*. The CDC then defines a pandemic as *an epidemic that has spread over several countries or continents, usually affecting a large number of people*[5]. Following these guidelines, COVID-19 was officially declared a pandemic on March 11th, 2020 by the World Health Organization (WHO) and is still considered to be an ongoing pandemic. Prior to COVID-19, the most recent public health crisis was Zika Virus, first isolated in Uganda in 1947 [6]. Other recent major infectious disease events include outbreaks of Ebola in 2014, MERS in 2013, SARS in 2003, and H1N1 Influenza (Swine Flu) in 2009.

As the world is battling the COVID-19 pandemic, one of the most vulnerable groups for mental health problems are frontline workers such as nurses, doctors, and emergency room staff. The risks of being on the front lines of combating the COVID-19 pandemic are not well understood [7], with even less being known about how to ensure the workers remain mentally healthy. Post-SARS research suggests hospital administration and staff needs to recognize that the impact to the health of frontline workers and frequent changing of infectious disease policy can play major roles in detrimentally impacting frontline workers mental health [8]. It is known that survivors of infectious diseases have higher rates of Post-Traumatic Stress Disorder (PTSD) [9]. It is also known that negative outcomes such as anxiety, burnout, and depression have been reported after outbreaks, suggesting possible long term effects of being on the front line during health crisis [10].

In this paper, we utilize survey data obtained from the University of Michigan's Inter-university Consortium for Political and Social Research (ICPSR). The data was collected by Deirdre Conroy [11] from the University of Michigan Department of Psychiatry and Cathy Goldstein from University of Michigan Department of Neurology. According to the ICPSR, "The rationale for this study was to examine whether sleep, mood, and health related behaviors might differ between healthcare workers

who transitioned to conducting care from home and those who continued to report in-person to their respective hospitals or healthcare facilities.” The original data contained 916 survey responses. The average survey response answered 94.5% of the survey questions, totalling 29 questions in total, with many questions leaving room for respondents to write how their mood or habits had changed since COVID-19 protocols were in effect. The data was stripped of any identifying information about the respondents and contained both categorical and numeric columns [12]. We would prefer not to use machine learning techniques or models (e.g. KNN) to treat the missing values for the categorical variables to keep the results reliable. We directly remove the datapoints with too many missing values from the dataset, and as the result, we end up with 518 data points.

Since the mental health dataset contains categorical and ordinal variables, we first encode them to numbers, and then if it is necessary, we use several encoding techniques such as one-hot encoding or dummy variable encoding as well as several packages in Python such as *OneHotEncoder*, *LabelEncoder* or *OrdinalEncoder* from *sklearn.preprocessing* to prepare the dataset for analysis.

Upon encoding all of the survey data, we began applying slightly more and more complicated models to the data, all being used to predict Question 29 (a), which reads “Please tell us how your mood has changed?” (see Figure 1). Then, by examining the models utilizing their feature importance, we identify the top predictors of mental health decline. The many different models applied to the data are consistent in their findings, identifying five key questions or groups of questions that are highly predictive of mental health decline in frontline workers.

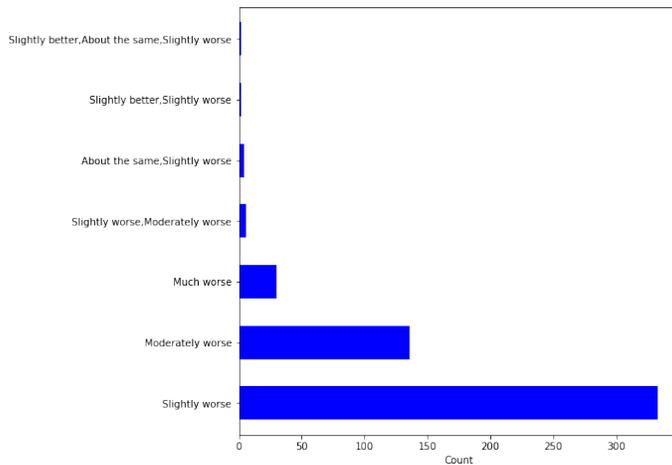


FIGURE 1 Q29 (a): “Please tell us how your mood has changed. My mood has been:”

2. Methodology

Our main goal in this paper is to use multiple supervised and unsupervised machine learning models and techniques to find the top predictors (features) of mental health data. Feature importance can be interpreted as the features that were the most valuable for generating the final prediction. Phrased differently, a high feature importance means the feature contains a lot of predictive power.

In this section, we provide an overview and results for all supervised and unsupervised machine learning methods that we use in this paper. To recall, the goal of supervised learning is to predict the value or the class of an unseen data point. To achieve this goal, we train a model on a training data and evaluate its accuracy on a test data. In all supervised methods, we split the data set into two parts, one part (for instance 75% of all observations) as the training data set, and the other part (for instance 25% of all data set) as the test data set. We train each model only on the training data set and then test the model on the test data set.

Each supervised predictor contains some hyper-parameters that control overfitting or underfitting. For instance, hyper-parameters for decision trees, k-nearest neighbors, neural networks are the depth of the tree, k, and the number of hidden layers, respectively. There are different techniques, such as the elbow rule, for finding the best value for the hyper-parameters [13]. Now, we overview the most common unsupervised and supervised techniques.

2.1. Unsupervised Feature Selection

Since the problem is a classification problem, where the majority of variables are categorical, we can use statistical tests such as Chi-Squared test to determine whether the target variable, Question 29 (a), is dependent or independent of the rest of variables. The variables that are independent can be considered as candidates for irrelevant features to the problem and they might be removed. During the data preparation process, we collapsed continuous variables into smaller groups (categories) to prepare the dataset for applying a Chi-Squared test [14]. After the data preparation process, we end up with a sample size of 513, which vastly exceeds the minimum of 20 to 50 recommended by Rana and Singhal [15] for avoiding Type II bias (failing to reject the null hypothesis when it is truly false).

We first construct a contingency table, calculate the expected frequencies for pairs of target variables and each input variable, and then apply a Chi-Squared test with significance level $\alpha = 0.05$

to determine whether or not there exists a correlation between frontline workers mental health decline and the rest of factors (features). It turns out that the Chi-Squared test rejects the following null hypotheses:

$$\begin{cases} H_0 : \text{Question 29 (a) is independent of Question } i, \\ H_1 : \text{Question 29 (a) is not independent of Question } i, \end{cases}$$

for $i \in \{13, 21, 22, 28\}$ (see Figures 2 and 3).

Chi-Squared test rejects the null hypothesis that the target variable is independent of the following variables
Q13. Approximately how many hours did you sleep on an average work night in the last week?
Q21. In January 2020, approximately how often did you use marijuana/cannabis (recreational or medical)?
Q22. In the last month, approximately how often did you use marijuana/cannabis (recreational or medical)?
Q28. Has the amount of food you have been eating per day changed?

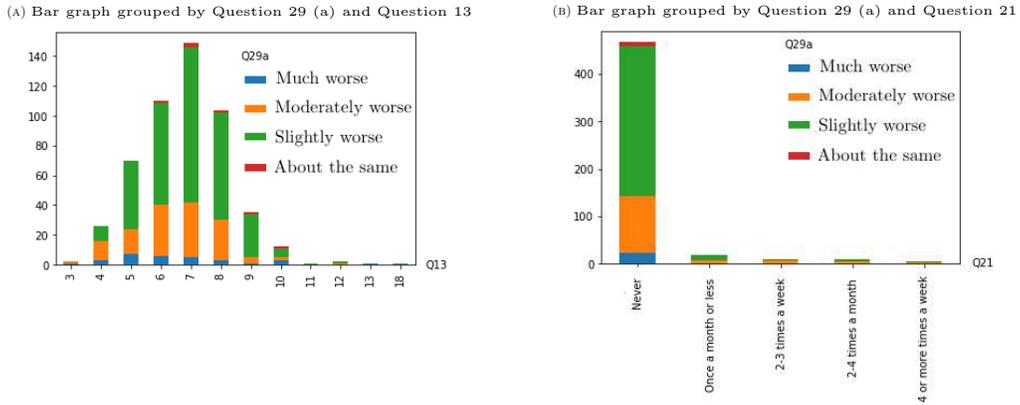


FIGURE 2 Bar graph grouped for the variables that pass the Chi-Squared test with significance level $\alpha = 0.05$

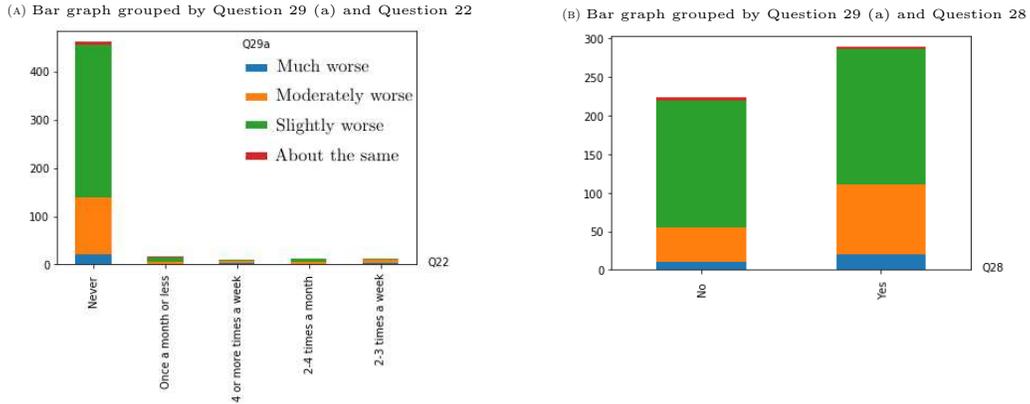


FIGURE 3 Bar graph grouped for the variables that pass the Chi-Squared test with significance level $\alpha = 0.05$

2.2. Supervised Feature Selection

2.2.1. Decision trees, Multinomial Logistic Regression and Recursive Feature Elimination

Decision trees are one of the most popular machine learning algorithms due to visualization simplicity. In other words, when we construct a tree, we can figure out what feature is most relevant at the first glimpse. A decision tree breaks down a dataset into smaller and smaller subsets containing data points that are more homogenous. Given a dataset with many features, different kind of decision trees with different depth can be constructed for training the model. However, finding the optimal decision tree, which is the shortest tree that can predict the class label for any unseen data, is computationally expensive. A decision tree contains a root node, internal nodes and leaf or terminal nodes. In a decision tree each leaf is assigned for a class label. Moreover, the root node and internal nodes contain feature test conditions. The order of non-terminal roots plays a significant role to find the optimal tree. A decision tree is also can be used for feature selection because it works based on the relevance of each feature. The goal of a decision tree is to minimize the entropy of the current sample to the next subsets by splitting the sample appropriately. The entropy for a sample S with n classes is defined as

$$E(S) = - \sum_{i=1}^n p_i \log(p_i), \quad (1)$$

where p_i is the probability of the class i in the sample S [20]. The accuracy of a decision tree model on the Mental Health Data is 74.12%.

Logistic regression is used to model binary categorical outcomes because it is improper to use linear regression while the output is not numeric, and the error terms are not normally distributed. In linear regression, parameters are estimated by ordinary least squares (OLS), where the sum of squared deviations of the predicted values from observed values is minimized. However, OLS is not appropriate for logistic regression to find unbiased estimators with minimum variance. For logistic regression, maximum likelihood estimation (MLE) is widely used to estimate parameters. Multinomial logistic regression is an appropriate extension of binary logistic regression for multiclass classification problems. It uses maximum likelihood estimation (MLE) to evaluate the probability of a target variable’s classes [17]. The accuracy of a multinomial logistic regression model on the Mental Health Data is 64.71%.

The Recursive Feature Elimination (RFE) is a machine learning technique for finding the most relevant feature. It works by recursively removing attributes and building a model on those attributes that remain. If the accuracy of a model drops significantly in absence of a feature, it indicates that the feature is important. We apply the recursive feature elimination using multinomial logistic regression and decision tree in presence and absence of all features. Then we rank features according to the reduction of accuracy of a model in absence of a feature at a time (see Table 1).

RFE feature selection using Decision Tree		RFE feature selection using Multinomial Logistic Regression	
Rank	Feature	Rank	Feature
#1	Q21: In January 2020, approximately how often did you use marijuana/cannabis?	#1	Q18: Has the amount of alcohol you are consuming changed?
#2	Q24: How many hours of COVID-19 related news or social media are you consuming on average per day?	#2	Q14: Has the number of your worked hours per week changed?
#3	Q20: In the last month, approximately how often did you have a drink containing alcohol?	#3	Q16: Have your sleep patterns changed?
#4	Q23: Has the amount of news you are consuming increased since the end of Feb, 2020?	#4	Q20: In the last month, approximately how often did you have a drink containing alcohol?

TABLE 1 Applying RFE on a Decision Tree and a Multinomial Logistic Regression

2.2.2. Naive Bayes, k -Nearest Neighbors, Support Vector Machine, Neural Networks and Random Forest

The Naive Bayes Classifier works based on the Bayesian theorem and is efficient when the dimensionality of the inputs is high. Given a data set comprising N observations with their corresponding

target class labels and D features $X = (x_1, x_2, \dots, x_d)$, we want to find the posterior probability for the outcome of the event Ω among a set of possible outcomes $\{y_1, y_2, \dots, y_K\}$. In a more familiar language, X is the predictors and $\{y_1, y_2, \dots, y_K\}$ is the set of categorical levels present in the dependent variable. So,

$$p(\Omega = y_j | x_1, x_2, \dots, x_d) \propto p(x_1, x_2, \dots, x_d | y_j) p(y_j) \quad (2)$$

The Naive Bayes classifier is assumed to be applied for datasets that the conditional probabilities of the features are statistically independent. So

$$p(\Omega = y_j | x_1, x_2, \dots, x_d) \propto p(x_1 | y_j) p(x_2 | y_j) \dots p(x_d | y_j) p(y_j). \quad (3)$$

The Naive Bayes classifier reduces a high-dimensional task to a one-dimensional kernel density estimation. It is an effective and commonly used probabilistic machine learning classifier. Naive Bayes classifiers are especially used for text classification and spam detection [21]. There is very little training in Naive Bayes compared to other common classification methods. The accuracy of a Naive Bayes classifier model on the Mental Health Data is 43.52%.

The k-nearest neighbors algorithm (KNN) classifies a new observation by a majority vote of its neighbors and it does not have the training phase. It calculates the distance from the new observation point to all seen data points, and then the new observation class label is assigned to be the class that is the most common among its k nearest neighbors. KNN should be one of the first choices for a classification when there is little or no prior knowledge about the data and feature labels. KNN is computationally expensive, especially when the dimension of feature vector is high, because it must store all data points and their distances from the new observation [22]. The accuracy of the KKN model on the Mental Health Data is 71.76%.

The support vector machine algorithm seeks a hyperplane in an n -dimensional space (n is the number of features) that classifies the data points with the maximum margin. The data points that are closer to the hyperplane and play an important role in finding the best position and orientation of the hyperplane to maximize the margine are called vector machines. One-vs-one Support Vector Machines (SVM OVO) is an appropriate extension of binary for multiclass classification problems. It splits the dataset into one dataset for each class versus every other class, which means it converts a multiclass classification dataset into multiple binary classification problems [18]. The accuracy of the SVM OVO model on the Mental Health Data is 70%.

Neural networks are a set of algorithms that are designed to recognize patterns. We can consider them as clustering and classification layers on top of the data we store. They predict the label of unseen data according to similarities among the example inputs. Neural networks can also be used for extracting features that are fed to other algorithms for classification. Deep learning is the name that is used for stacked neural networks, which contains several layers. Each layer in the network is made of different neurons, where computations happen to decide whether a neuron should fire and the signal should progress further through the network. A neuron fires when it meets a sufficient stimuli. A neuron combines the last layer neurons output with a set of weights. This weighted sum is the input of a function called activation function to determine whether and to what extent the signal should progress to affect the ultimate outcome [19]. The accuracy of a neural network with 10 hidden layers and softmax as the last layer activation function on the Mental Health Data is 82.35%.

Random Forest Classifier is an ensemble algorithm, which is a model that combines more than one algorithm of same or different kind for classifying objects. From randomly selected subset of training data, random forest generates a set of decision trees, and then it specifies the class label of a data point by aggregating the votes from different decision trees. The random forest is more powerful than a single decision tree classifier because it avoids overfitting on the training data. The accuracy of a regular random forest with 10 trees on the Mental Health Data is 80.59%.

The accuracy scores of the above models together is shown in Figure 4. Moreover, Figure 5 displays feature scores of Random Forest Model.

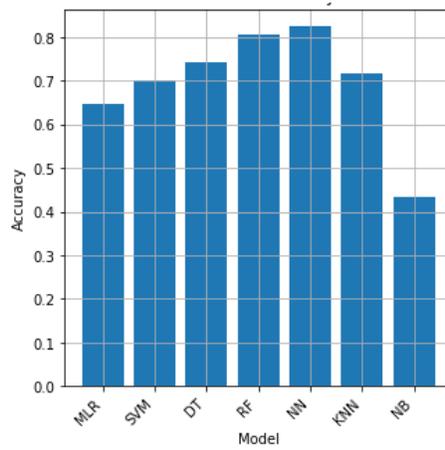


FIGURE 4 Machine learning models v.s. accuracy

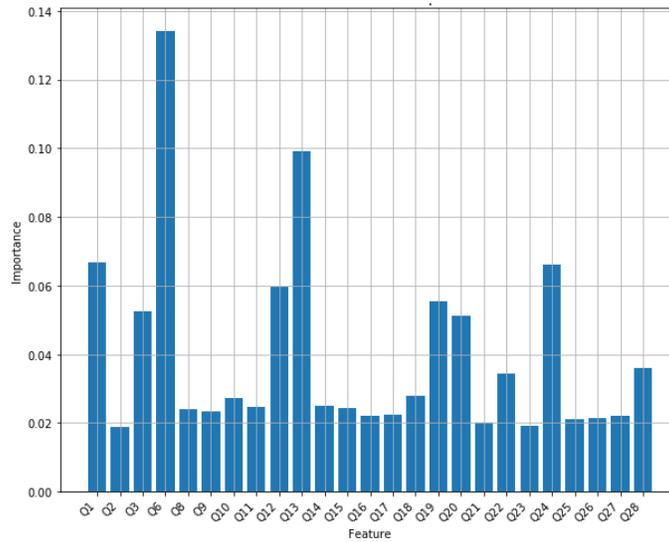


FIGURE 5 Feature importance scores of Random Forest

2.2.3. Gradient Tree Boosting, XGBoost, CatBoost, LightGBM and Feature Importances

Gradient tree boosting is an ensemble model, which is constructed from decision tree models. Decision trees are added one at a time to the model and fit to minimize the error made by previous trees. In this model, gradient descent or stochastic gradient descent methods are used to minimize the differentiable loss functions. Gradient boosting is one of the most widely used machine learning

algorithm on tabular datasets.

The number of decision trees is one of the most important hyperparameter for the Gradient Boosting ensemble algorithm. The depth of the trees and the number of trees can be used efficiently in the ensemble to improve the performance of a model. Figure 6 displays a box and whisker plot and the effect of the number of trees on accuracy. It turns out that if we change the number of trees between 16 and 19, the mean accuracy scores of ensemble models do not change significantly. The mean accuracy of the ensemble model meets its maximum 87.1% when the number of trees is between 16 and 19. Figure 7 displays the feature scores Gradient Tree Boosting with 19 trees.

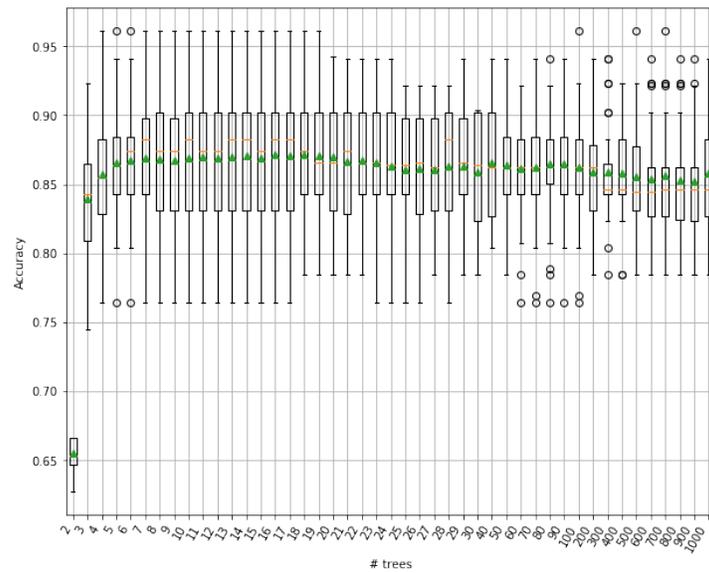


FIGURE 6 Box plot for the mean accuracy of Gradient Tree Boosting with different number of trees

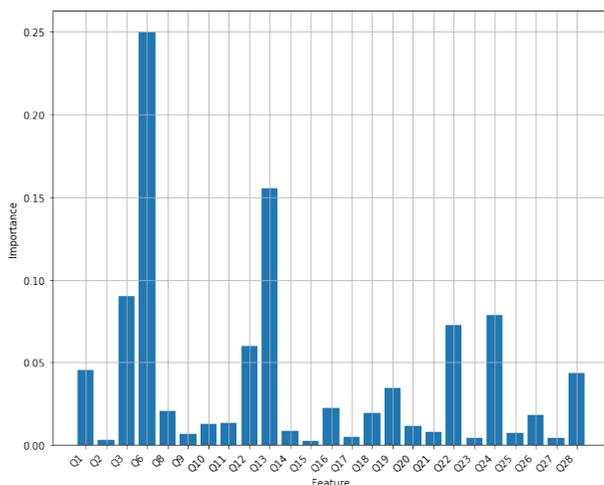


FIGURE 7 The feature scores of Gradient Tree Boosting with 19 trees

Gradient boosted trees have become the most widely used machine learning algorithms when it comes to tabular data. There are several popular boosting algorithms such as XGBoost, CatBoost and LightGBM. However, XGBoost, CatBoost and LightGBM algorithms differ from one another in the implementation of the boosted trees algorithm and the splitting mechanism.

XGBoost, introduced by Tianqi Chen [23, 24], is a widely used machine learning algorithm that implements machine learning algorithms using the Gradient Boosting framework with high accuracy and speed. It was initially introduced to improve GBM’s training time. It provides a parallel tree boosting framework to solve so many multiclass classification problems. Cross-Validation is a popular method to find a better accuracy for a model rather than a simple train/test split. The cross validation procedure is: shuffling the data set randomly, splitting the data set in k groups, taking each group as a test data and feeding the rest to learn a model and find its accuracy, and then calculating and returning the mean of all k obtained accuracy. If the value for k is assigned to be the number of observations, then it is called leave-one-out cross-validation. Figure 8 displays the accuracy scores of XGBoost with multiple folds ($2 \leq k \leq 20$) for cross-validation. It turns out that XGBoost with $k=9$ folds gives the best accuracy, approximately 86%. Figure 9 displays the feature scores of the XGBoost model using $k = 9$ cross-validation.

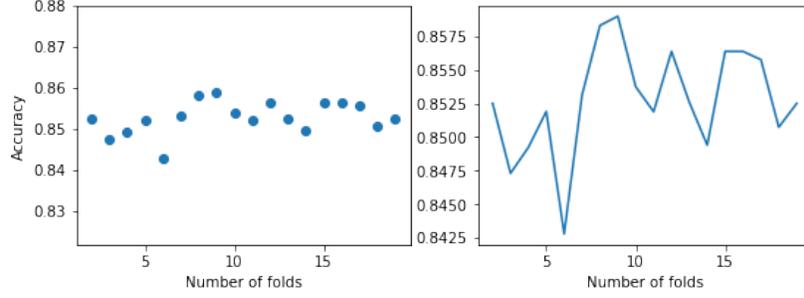


FIGURE 8 The mean accuracy of XGBoost using k-fold cross-validation

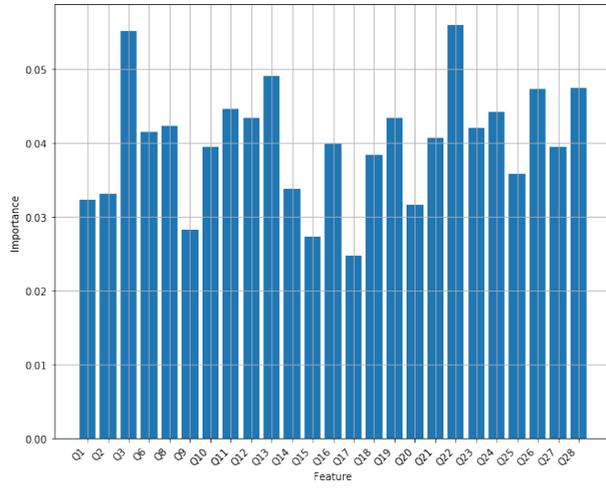


FIGURE 9 The feature scores of XGBoost using $k = 9$ cross-validation

CatBoost [25] is another algorithm for gradient boosting on decision trees. CatBoost trains trees sequentially, so that each successive tree is built with reduced loss compared to the previous trees. CatBoost can be used directly on categorical variables without any explicit pre-processing to convert categories into numbers. The numbers of trees is one of the most important hyper-parameters in CatBoost algorithm. Figure 10 displays the accuracy of CatBoost model on the Mental Health dataset for multiple number of trees. It turns out that CatBoost with 60 trees gives the best accuracy, approximately 85.50%. Figure 11 displays the feature scores of the CatBoost model using 60 trees.

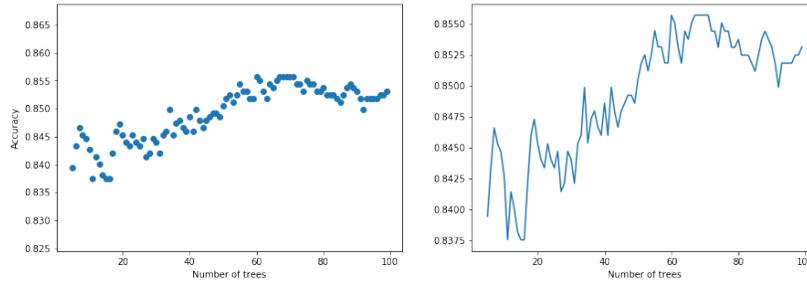


FIGURE 10 The mean accuracy of CatBoost using several trees

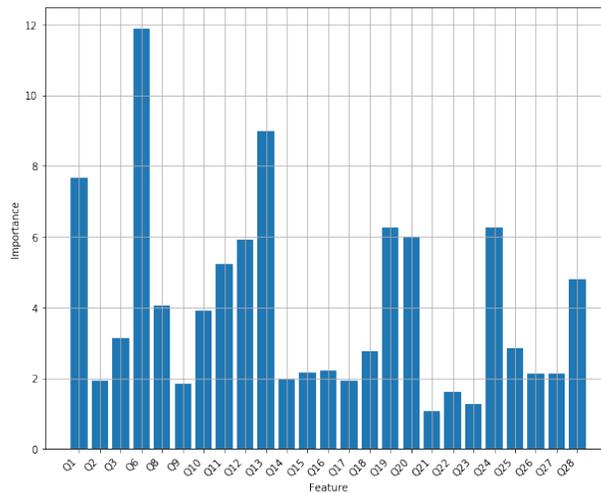


FIGURE 11 The feature scores of the CatBoost model using 60 trees

LightGBM, introduced by Guolin Ke et. al. [26], is another strong learner that applies the boosting framework, which train trees sequentially instead of simultaneously, like CatBoost. LightGBM uses gradient-based one-side sampling (GOSS) that aims to find a good balance between increasing speed by reducing the number of data points and holding the accuracy for learned trees. LightGBM differs from XGBoost and CatBoost in the way that it works based on gradient-based one-side sampling (GOSS) and Exclusive Feature Bundling (EFB).

GOSS excludes a significant proportion of data instances with small gradients, using the rest to estimate the information gain in individual trees. Guolin Ke et. al. [26] prove the data instances with larger gradients play a more important role in the computation of information gain, and thus GOSS can obtain an accurate estimation of the information gain with a much smaller data size. EFB

bundles mutually exclusive features to reduce the number of features. Mutually exclusive features are features that rarely take nonzero values simultaneously, such as one-hot encoded features. Guolin Ke et. al. [26] prove that finding the optimal bundling of exclusive features is NP-hard, but a greedy algorithm can achieve a good approximation ratio, effectively reducing the number of features without hurting accuracy of the model.

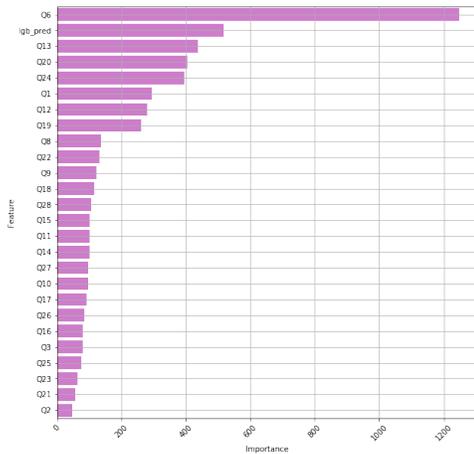
Tables 2 and 3 display the accuracy scores of LightGBM model stacked once and twice, respectively. The model uses 200 gradient-boosted decision trees (GBDT) all limited to a maximum depth of 8. Figure 12 displays the feature scores of LightGBM model stacked with one and two variables on the Mental Health data. The first four important features for LightGBM stacked with one and two variables are shown in Table 4.

	precision	recall	f1-score	support
accuracy			0.81	154
macro avg	0.46	0.44	0.43	154
weighted avg	0.77	0.81	0.78	154

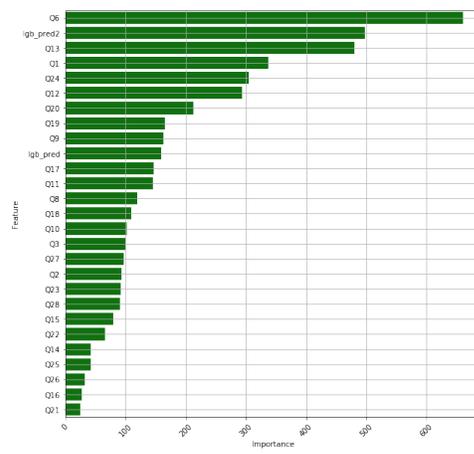
TABLE 2 Stacked LightGBM Scores

	precision	recall	f1-score	support
accuracy			0.91	180
macro avg	0.46	0.46	0.46	180
weighted avg	0.90	0.91	0.90	180

TABLE 3 2x Stacked LightGBM Scores



(A) LightGBM stacked x1 variable importance



(B) LightGBM stacked x2 variable importance

FIGURE 12 LightGBM

LightGBM stacked with 1 variables		LightGBM stacked with 2 variables	
Rank	Feature	Rank	Feature
#1	Q6: What is your role in the healthcare field? (e.g. psychologist, physician, nurse)	#1	Q6: What is your role in the healthcare field? (e.g. psychologist, physician, nurse)
#2	Q13: Approximately how many hours did you sleep on an average work night in the last week?	#2	Q13: Approximately how many hours did you sleep on an average work night in the last week?
#3	Q20: In the last month, approximately how often did you have a drink containing alcohol?	#3	Q1: What is your age?
#4	Q24: How many hours of COVID-19 related news or social media are you consuming on average per day?	#4	Q24: How many hours of COVID-19 related news or social media are you consuming on average per day?

TABLE 4 The most important variables for LightGBM

2.2.4. Synthetic Minority Oversampling Technique

Note that the categories of the target variable, Question 29 (a), are not approximately equally represented. Nitesh Chawla, et al. [16] proposed a technique called the Synthetic Minority Oversampling Technique, or SMOTE, for synthesizing new examples of the minority classes. The SMOTE is a combination of over-sampling the minority classes and under-sampling the majority classes to achieve a better classifier performance. Since there are few examples with labels belong to some classes, we remove them before we apply SMOTE to avoid misleading the model (see Figure 15a).

With SMOTE, we can customize the oversampling strategy. Figure 13 displays the first customized oversampling strategy and the feature importance of the model with its accuracy.

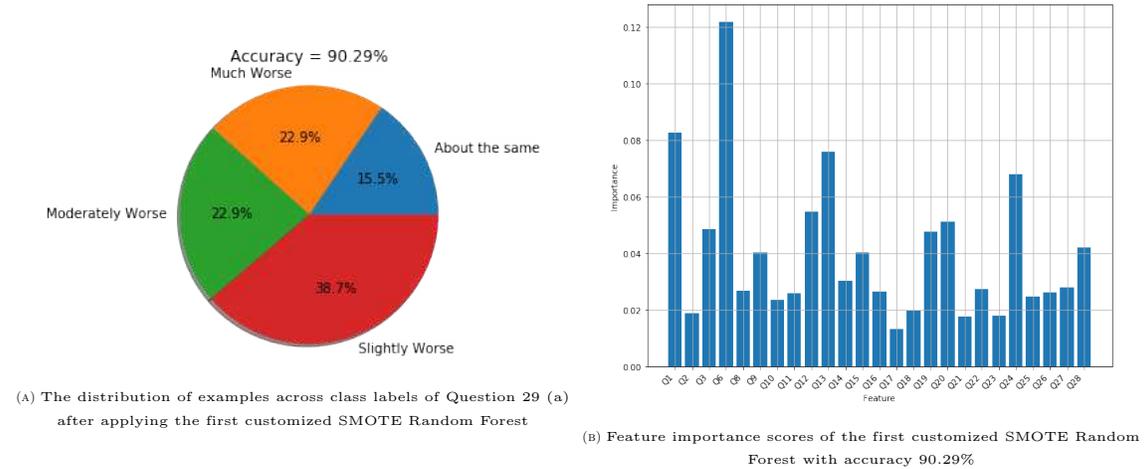


FIGURE 13 The distribution of examples across class labels of Question 29 (a), feature importance scores and the obtained accuracy score

Figure 14 displays the second and the third customized strategies for oversampling and the obtained accuracy scores for the models.

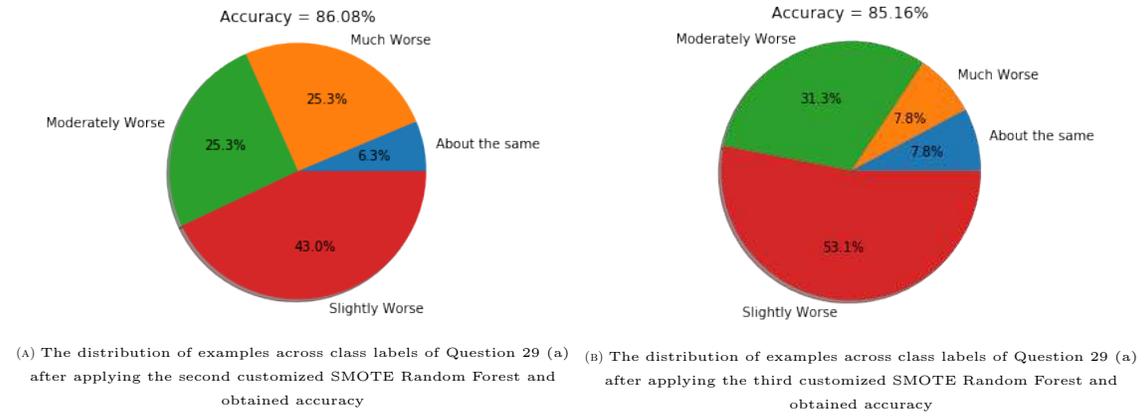


FIGURE 14 The distribution of examples across class labels of Question 29 (a) and the obtained accuracy scores

However, applying SMOTE to oversample all classes to the number of examples in the majority class on Random Forest improves accuracy of the model from 80.59% to 96.32% (see Figure 16) with questions 6, 13, 1 and 24 as the most important features, respectively (see Figure 17). The first four important features for the regular SMOTE Random Forest are shown in Table 5.

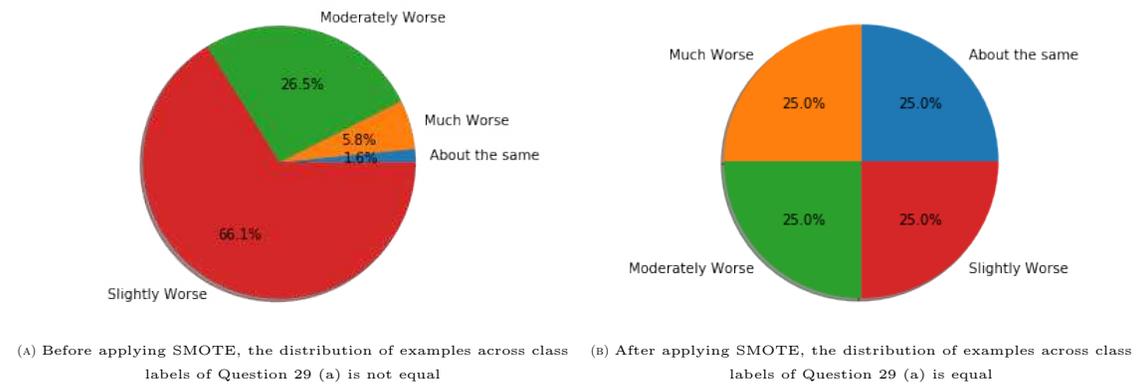


FIGURE 15 Applying SMOTE to oversample all classes to the number of examples in the majority class

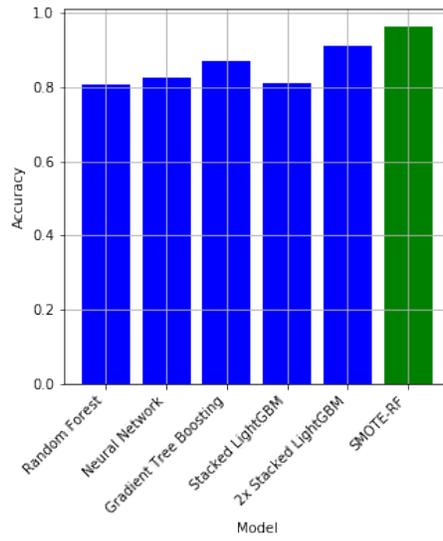


FIGURE 16 Improving the Random Forest accuracy by means of SMOTE

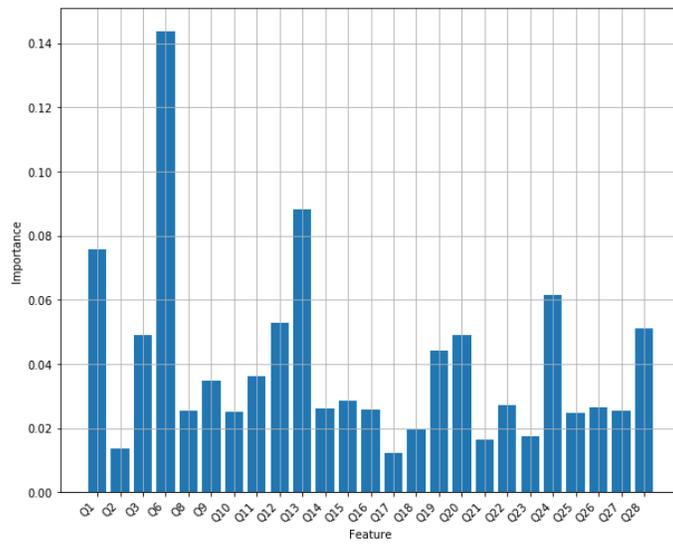


FIGURE 17 Feature importance scores of SMOTE Random Forest

SMOTE Random Forest	
Rank	Feature
#1	Q6: What is your role in the healthcare field? (e.g. psychologist, physician, nurse)
#2	Q13: Approximately how many hours did you sleep on an average work night in the last week?
#3	Q1: What is your age?
#4	Q24: How many hours of COVID-19 related news or social media are you consuming on average per day?

TABLE 5 The first four most important features for the regular SMOTE Random Forest

3. Discussion

Upon the examination of the top predictors from the many different models, we have identified four main questions as well as a group of questions (questions 18-22) as highly predictive of mental health decline in frontline workers. With the emergence of new strains of COVID-19 such as the Delta and Omicron variants, we may be able to learn from our past and improve the working environment for frontline workers such that their mental health does not decline much, if at all.

Our analysis concludes that Question 6: “What is your role in the healthcare field?” is the most important predictor of mental health decline. Both stacked LightGBM models, the Random Forest with and without SMOTE, CatBoost, and Gradient Tree Boosting all conclude that this is the most important predictor. This result is unsurprising, as we do not expect a frontline worker dealing with extremely sick COVID-19 patients would suffer a greater mental health decline than a child pediatrician who is not interacting with COVID-19 patients. One could find in the face of the recent variants, such as the Omicron variant, that these issues may flare up again, and there may be no one size fits all solution for the medical institution seeking to preserve the mental health of its employees across different disciplines. As such, implementing role specific support groups may be able to relieve some of the stress different workers may be facing, as sharing experiences with someone who has gone through or is currently experiencing the same things may allow for the workers to bond over their adversities and preserve their mental health.

We also see a clear second most important predictor in Question 13: “Approximately how

many hours did you sleep on an average work night in the last week?” The Random Forest with and without SMOTE, Gradient Tree Boosting, CatBoost, both LightGBM models, and the Chi Squared analysis find this is a highly important predictor. We know that COVID-19 if contracted can have large, potentially deadly, effects on frontline workers. However, this analysis reveals one of the “hidden” effects COVID-19 has as well. Even without contracting the virus, frontline workers are having impacted sleeping schedules, as the stress of their jobs takes a toll on their body both mentally and physically. This may also be part of the institutions they are employed by asking for longer hours to compensate for a lack of frontline workers or an excess of patients sick with the virus. The institutions may find that the frontline workers’ performance and level of care may increase if they resume their normal sleep schedules as well as the mental health of their workers improving.

Our analysis finds that Question 24: “How many hours of COVID-19 related news or social media are you consuming on average per day?” is a highly important predictor of mental health decline in frontline workers. The Random Forest with and without SMOTE, Gradient Tree Boosting, both LightGBM models, and the Recursive Feature Elimination utilizing a decision tree all find this is in the top 5 most important predictors. At first, we found this result surprising, but after further deliberation, we believe that this may be the easiest factor to improve on. There is seldom positive news regarding the virus circulating on the news or on social media, with many headlines pointing to new variants that are more contagious or deadly than the one before it, or how we are breaking new daily case records, and even that other states and countries are re-entering lockdowns. It is not surprising to see that consuming more and more of this media can lead one to doubt that we will ever beat COVID-19 and that the fight may be futile, leading to a decrease in mental health. One could purposefully try to limit their consumption of such media and may find that their mental health and possibly even physical health (such as sleep patterns) may improve.

Looking at the top predictors of the Random Forest with and without SMOTE, CatBoost, and both LightGBM models, we find that Question 1: “What is your age?” is one of the top five most important predictors of mental health decline in frontline workers. We hypothesize that workers that are younger and less familiar with intense medical situations may be more prone to experiencing mental health decline upon encountering said situations. Healthcare institutions may want to encourage their less and more seasoned staff to work together, allowing for the younger staff to learn helpful tips for their job as well as enabling them to reach out to the seasoned staff if

they feel as though they are beginning to have mental health issues.

Lastly, we find that a group of questions regarding alcohol and cannabis use stand out as highly predictive in many different models such as the Recursive Feature Elimination with both Decision Trees and Multinomial Regression, the Chi Squared test, the Random Forest with and without SMOTE, XGBoost, CatBoost, and both LightGBM models. This result is intuitive, as many people may turn to different methods of coping with the extreme stress that they may be feeling during their day-to-day lives combating COVID-19 and its many variants. Healthcare institutions may find that offering counseling services and other stress relief programs may decrease the usage of Alcohol and Cannabis in their frontline workers, improving both their mental and physical health.

Our analysis has demonstrated that COVID-19 has not taken a toll on the mental health of frontline workers, but their physical health too. With the emergence of many variants such as the Delta variant and recent Omicron variant, large healthcare institutions can learn from the past to improve conditions and offer their frontline workers a better chance of improving/maintaining their mental and physical health. We know some of the physical symptoms that come from contracting COVID-19, but what we have found is that there exist some “hidden” effects of the virus that do not necessarily come from direct contact. These “hidden” effects may be partially preventable, and in the face of new variants that could cause a resurgence of the virus, healthcare institutions must take steps to prevent these “hidden” symptoms from becoming problems for our frontline workers. Some possible remedies may be easier to implement than others, such as warning staff about consuming too much COVID-19 related news and offering counseling. Even working to allow younger staff to work with older, more experienced staff may result in better training for the institution while simultaneously working to better their staff’s combined mental health.

4. Conclusion and Future Work

We have presented an analysis of a COVID-19 mental health survey data obtained from the University of Michigan Inter-University Consortium for Political and Social Research. As part of the analysis, we have utilized a variety of statistical and machine learning models and techniques such as Decision Trees, Multinomial Logistic Regression, Naive Bayes, k-Nearest Neighbors, Support Vector Machines, Neural Networks, Random Forests, Gradient Tree Boosting, XGBoost, CatBoost, LightGBM, Synthetic Minority Oversampling, and a Chi Squared Test. Through the interpretation of the many models applied to the mental health survey data, we have concluded that the most

important factor in predicting the mental health decline of a frontline worker is the healthcare role the individual is in (Nurse, Emergency Room Staff, Surgeon, etc.), followed by the amount of sleep the individual has had in the last week, the amount of COVID-19 related news an individual has consumed on average in a day, the age of the worker, and the usage of alcohol and cannabis. In light of the recent identification of the Omicron and Delta variants of COVID-19, we hope that these findings can be utilized by healthcare facilities to help preserve or improve their employee's mental health.

In future work we would like to aggregate more data on frontline workers, their habits, and their mental health. With more data from a diverse range of locations, we would have the ability to apply even more complex and accurate models to the data, while simultaneously allowing us to make even stronger conclusions about the impacts COVID-19 has on the mental health of frontline workers. Additionally, in future work, we may be able to utilize permutation importance along with the feature importance to shed even more light on the inner workings of the matching learning models. Finally, we would like to identify specific features then utilize generalized linear models to more exactly quantify the relationship the variables have with mental health decline.

We are also interested in analyzing accuracy, speed and feature scores of some machine learning models on the COVID-19 mental health data when we replace the most common optimizers such as GD, SGD or Limited-memory BFGS with the ones that are introduced in [27] and [28].

5. Declarations

- Funding: Not Applicable
- Conflicts of interest/Competing interests: Not Applicable
- Ethics approval: Not Applicable
- Consent to participate: Not Applicable
- Consent for publication: Not Applicable
- Availability of data and material: It is Available: <https://www.openicpsr.org/openicpsr/project/127081/version/V1/view?path=/openicpsr/127081/fcr:versions/V1&type=project>

- Code availability: It is Available. The codes will be available online once the paper is accepted.
- Authors' contributions: We analyze a COVID-19 mental health survey data obtained from the University of Michigan Inter-University Consortium for Political and Social Research. As part of the analysis, we use a variety of statistical and machine learning models and techniques such as Decision Trees, Multinomial Logistic Regression, Naive Bayes, k-Nearest Neighbors, Support Vector Machines, Neural Networks, Random Forests, Gradient Tree Boosting, XG-Boost, CatBoost, LightGBM, Synthetic Minority Oversampling, and a Chi Squared Test. We conclude that the most important factor in predicting the mental health decline of a frontline worker is the healthcare role the individual is in (Nurse, Emergency Room Staff, Surgeon, etc.), followed by the amount of sleep the individual has had in the last week, the amount of COVID-19 related news an individual has consumed on average in a day, the age of the worker, and the usage of alcohol and cannabis. In light of the recent identification of the Omicron and Delta variants of COVID-19, we hope that these findings can be utilized by healthcare facilities to help preserve or improve their employee's mental health.

References

- [1] Huang, Chaolin, et al. "Clinical features of patients infected with 2019 novel coronavirus in Wuhan, China." *The lancet* 395.10223 (2020): 497-506.
- [2] Nishiura, Hiroshi, et al. "The extent of transmission of novel coronavirus in Wuhan, China, 2020." (2020): 330.
- [3] Harcourt, Jennifer, et al. "Severe acute respiratory syndrome coronavirus 2 from patient with coronavirus disease, United States." *Emerging infectious diseases* 26.6 (2020): 1266.
- [4] WHO Coronavirus Disease (COVID-19) Dashboard. World Health Organization. 2020 Aug 31. URL: <https://covid19.who.int/> [accessed 2020-09-01]
- [5] Centers for Disease Control and Prevention. "Lesson 1: Introduction to Epidemiology, Section 11: Epidemic Disease Occurrence." CDC. Available online at: <https://www.cdc.gov/csels/dsepd/ss1978/lesson1/section11.html>. [Last accessed on 5th May, 2020] (2018).
- [6] Yadav, Sankalp, Gautam Rawal, and Mudit Baxi. "Zika virus: a pandemic in progress." *Journal of translational internal medicine* 4.1 (2016): 42.

- [7] Cabarkapa, Sonja, et al. "The psychological impact of COVID-19 and other viral epidemics on frontline healthcare workers and ways to address it: A rapid systematic review." *Brain, behavior, & immunity-health* (2020): 100144.
- [8] Tam, Cindy WC, et al. "Severe acute respiratory syndrome (SARS) in Hong Kong in 2003: stress and psychological impact among frontline healthcare workers." *Psychological medicine* 34.7 (2004): 1197-1204.
- [9] Hong, Xia, et al. "Posttraumatic stress disorder in convalescent severe acute respiratory syndrome patients: a 4-year follow-up study." *General hospital psychiatry* 31.6 (2009): 546-554.
- [10] Lancee, W. J., R. G. Maunder, and D. S. Goldbloom. "The co-authors of the Impact of SARS Study. The prevalence of mental disorders in Toronto hospital workers one to two years after SARS." *Psychiatr Serv* 59.1 (2008): 91-95.
- [11] Conroy, Deirdre, and Goldstein, Cathy. COVID Isolation on Sleep and Health in Healthcare Workers. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2020-11-20. <https://doi.org/10.3886/E127081V1>
- [12] <https://www.openicpsr.org/openicpsr/project/127081/version/V1/view?path=/openicpsr/127081/fcr:versions/V1&type=project>
- [13] Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. "Machine learning basics." *Deep learning* 1.7 (2016): 98-164.
- [14] McHugh, Mary L. "The chi-square test of independence." *Biochemia medica* 23.2 (2013): 143-149.
- [15] Rana, Rakesh, and Richa Singhal. "Chi-square test and its application in hypothesis testing." *Journal of the Practice of Cardiovascular Sciences* 1.1 (2015): 69.
- [16] Chawla, Nitesh V., et al. "SMOTE: synthetic minority over-sampling technique." *Journal of artificial intelligence research* 16 (2002): 321-357.
- [17] Böhning, Dankmar. "Multinomial logistic regression algorithm." *Annals of the institute of Statistical Mathematics* 44.1 (1992): 197-200.

- [18] Brereton, Richard G., and Gavin R. Lloyd. "Support vector machines for classification and regression." *Analyst* 135.2 (2010): 230-267.
- [19] Bishop, Christopher M. "Pattern recognition." *Machine learning* 128.9 (2006).
- [20] Kotsiantis, Sotiris B. "Decision trees: a recent overview." *Artificial Intelligence Review* 39.4 (2013): 261-283.
- [21] Domingos, Pedro, and Michael Pazzani. "On the optimality of the simple Bayesian classifier under zero-one loss." *Machine learning* 29.2 (1997): 103-130.
- [22] Jiang, Liangxiao, et al. "Survey of improving k-nearest-neighbor for classification." *Fourth international conference on fuzzy systems and knowledge discovery (FSKD 2007)*. Vol. 1. IEEE, 2007.
- [23] Chen, Tianqi, and Carlos Guestrin. "Xgboost: A scalable tree boosting system." *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. 2016.
- [24] Chen, Tianqi, et al. "Xgboost: extreme gradient boosting." *R package version 0.4-2 1.4* (2015): 1-4.
- [25] Dorogush, Anna Veronika, Vasily Ershov, and Andrey Gulin. "CatBoost: gradient boosting with categorical features support." *arXiv preprint arXiv:1810.11363* (2018).
- [26] Ke, Guolin, et al. "Lightgbm: A highly efficient gradient boosting decision tree." *Advances in neural information processing systems* 30 (2017): 3146-3154.
- [27] Rezapour, Mostafa, and Thomas J. Asaki. "Adaptive trust-region algorithms for unconstrained optimization." *Optimization Methods and Software* (2019): 1-23.
- [28] Erway, Jennifer B., and Mostafa Rezapour. "A New Multipoint Symmetric Secant Method with a Dense Initial Matrix." *arXiv preprint arXiv:2107.06321* (2021).

Appendix A. Some important features in the mental health data

Question	Levels	Question	Levels
Q1: What is your age?	10-20 21-30 31-40 41-50 51-60 61-70 71-80 81-90	Q6: What is your role in the healthcare field?	Psychologist Physician Nurse Psychiatrist Social Worker Research coordinator/manager/Public health/Staff etc.
Q9: Are you conducting clinical video visits with your patients from your home?	No Yes	Q10: Does your work require you to follow fixed work schedule while working from home?	No Yes
Q14: Has the number of your worked hours per week changed?	No Yes	Q15: Have you varied your work schedule?	No Yes
Q16: Have your sleep patterns changed?	No Yes	Q17: Has the number of naps you are taking changed?	No Yes
Q18: Has the amount of alcohol you are consuming changed?	No Yes	Q19: In January 2020, approximately how often did you have a drink containing alcohol?	4 or more times a week 2-4 times a month 2-3 times a week Once a month or less Never
Q20: In the last month, approximately how often did you have a drink containing alcohol?	4 or more times a week 2-4 times a month 2-3 times a week Once a month or less Never	Q21: In January 2020, approximately how often did you use marijuana/cannabis?	4 or more times a week 2-4 times a month 2-3 times a week Once a month or less Never

Question	Levels	Question	Levels
Q22: In the last month, approximately how often did you use marijuana/cannabis?	4 or more times a week 2-4 times a month 2-3 times a week Once a month or less Never	Q23: Has the amount of news you are consuming increased since the end of Feb, 2020?	No Yes
Q24: How many hours of COVID-19 related news or social media are you consuming on average per day?	4 or more times a week 2-4 times a month 2-3 times a week Once a month or less Never	Q25: Have you had more "screen time" e.g. use of smartphone, tablet, etc. around bedtime?	No Yes
Q26: Have you changed your movement/exercise?	No Yes	Q27: Has the quality of your diet changed?	No Yes
Q28: Has the amount of food you have been eating per day changed?	No Yes	Q29: Has your mood changed?	No Yes
Q8: Are you currently conducting your job mostly from home now?	No Yes	Q29 (a): Please tell us how your mood has changed. My mood has been:	Much better Moderately better Slightly better About the same Slightly worse Moderately worse Much worse

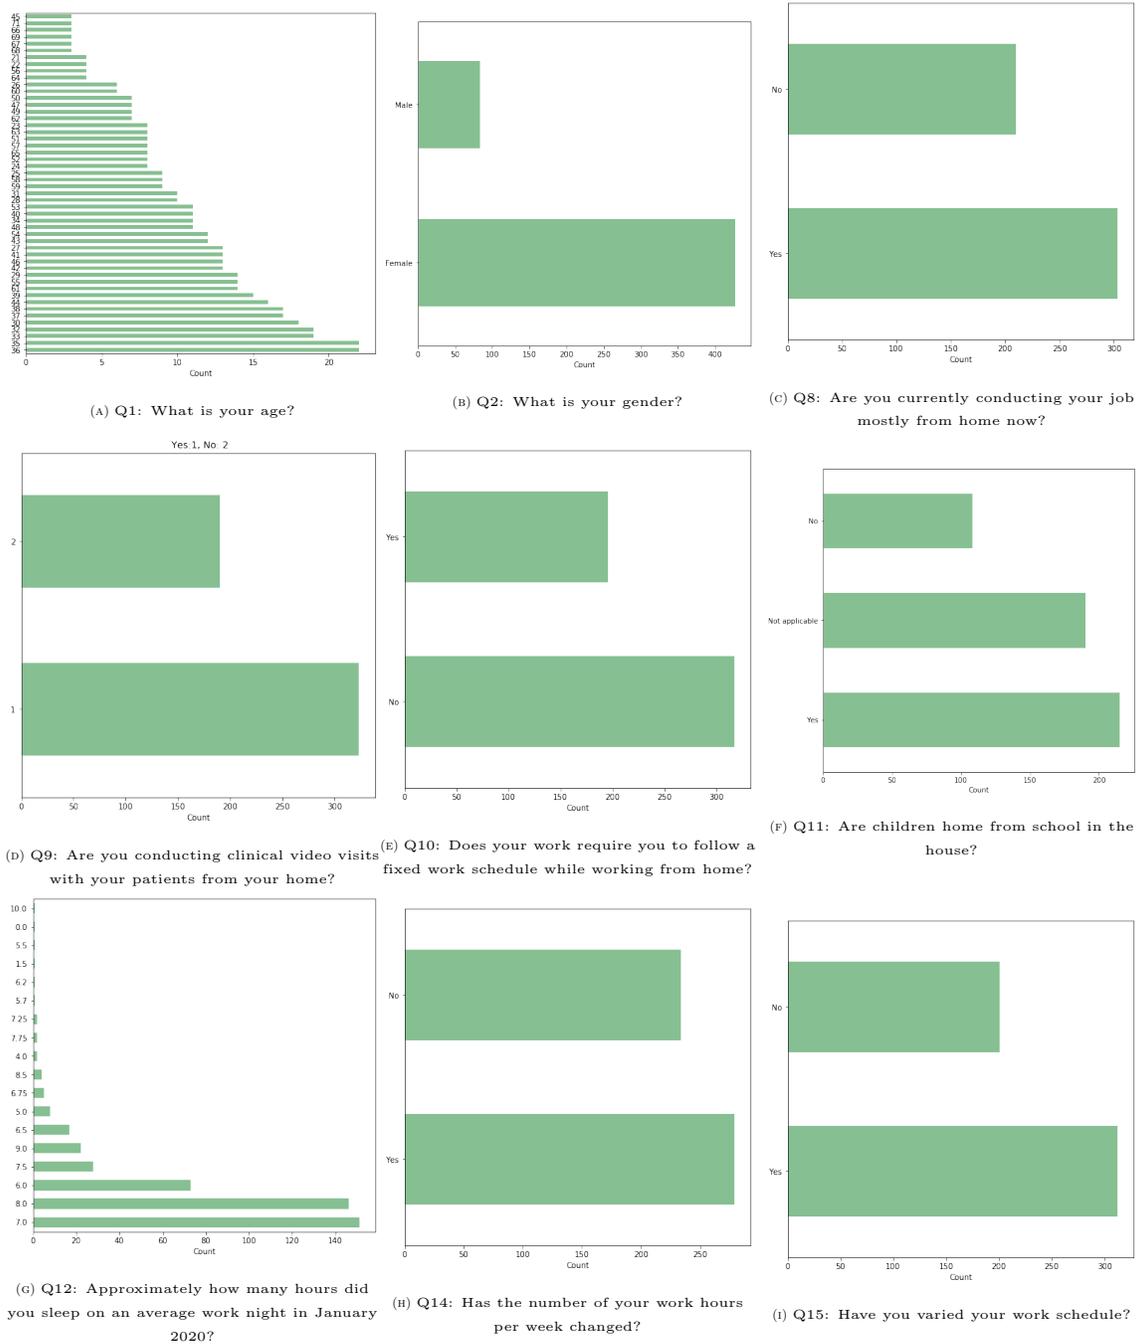


FIGURE A.18 Frequency bar charts after the preprocessing phase

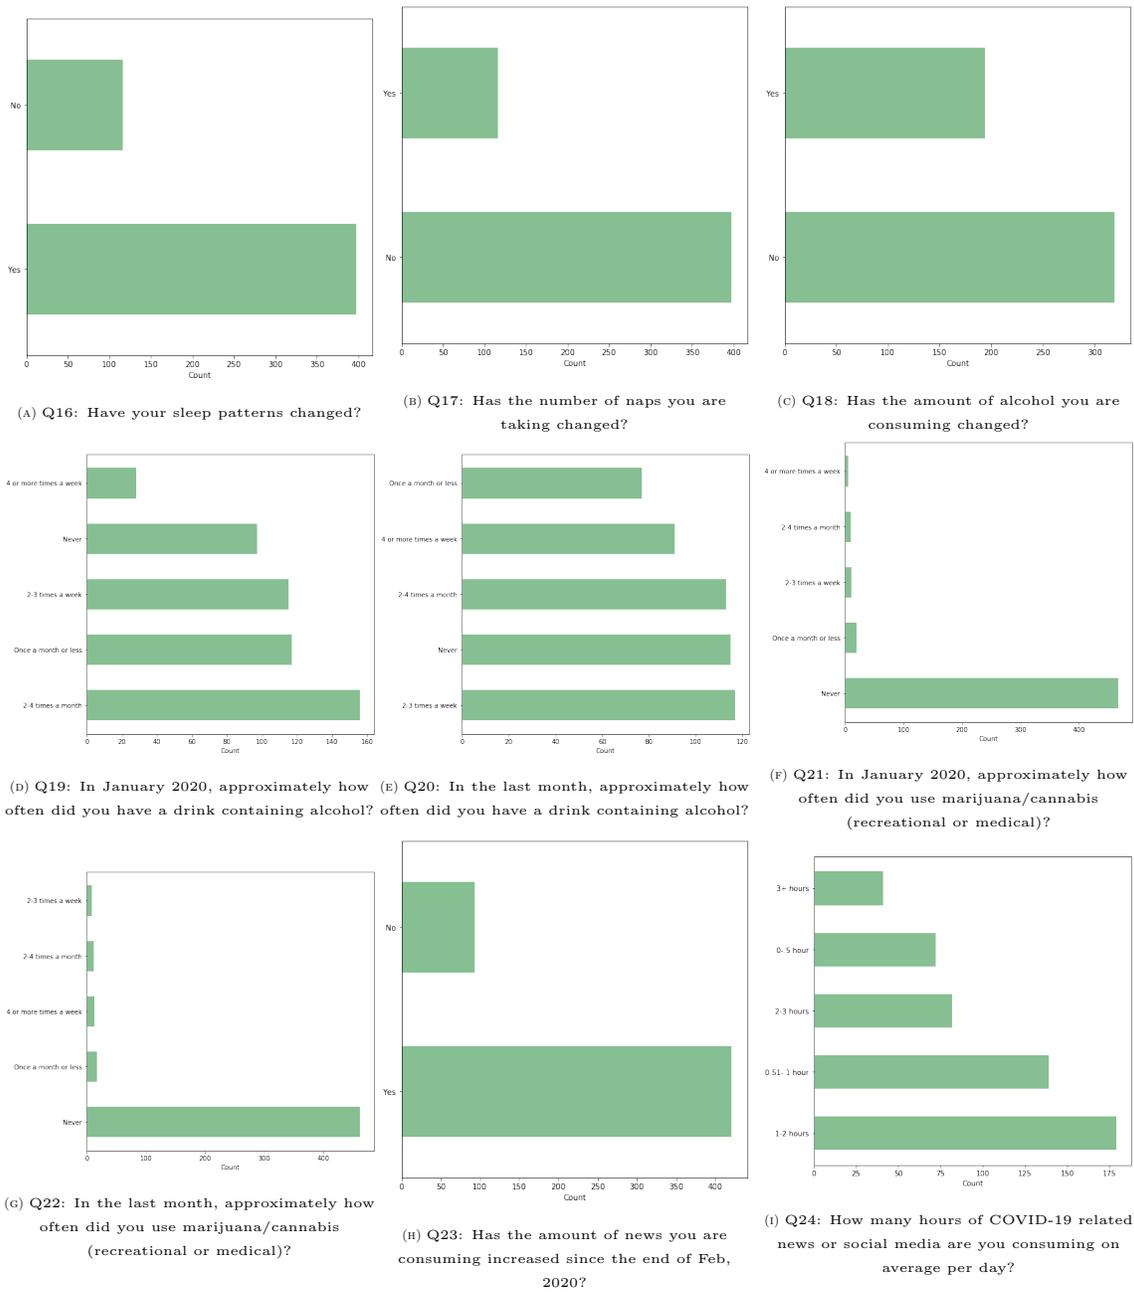
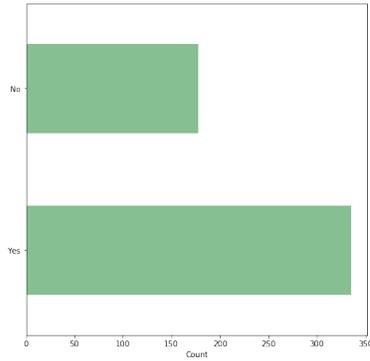
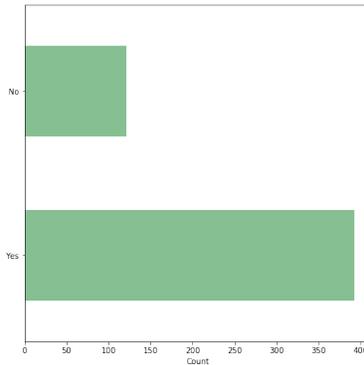


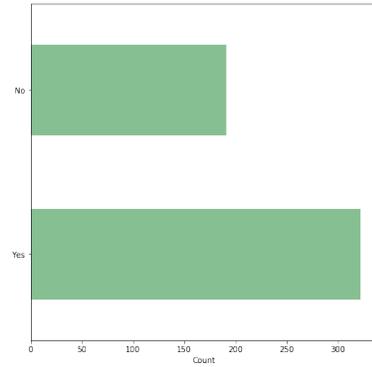
FIGURE A.19 Frequency bar charts after the preprocessing phase



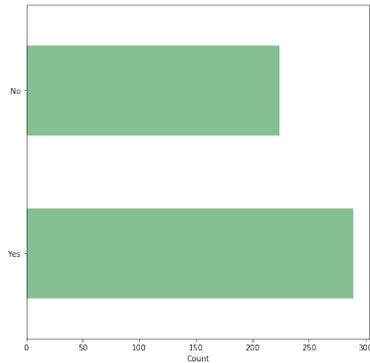
(A) Q25: Have you had more “screen time” (e.g. use of smartphone, tablet, etc.) around bedtime?



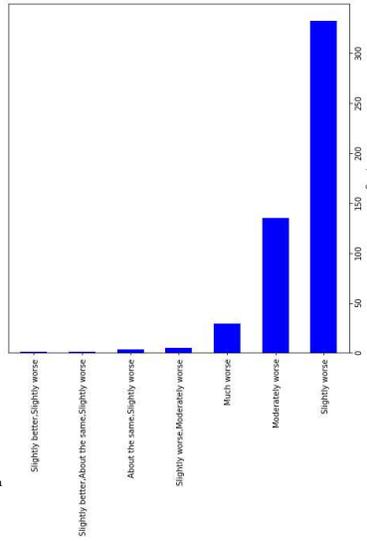
(B) Q26: Have you changed your movement/exercise?



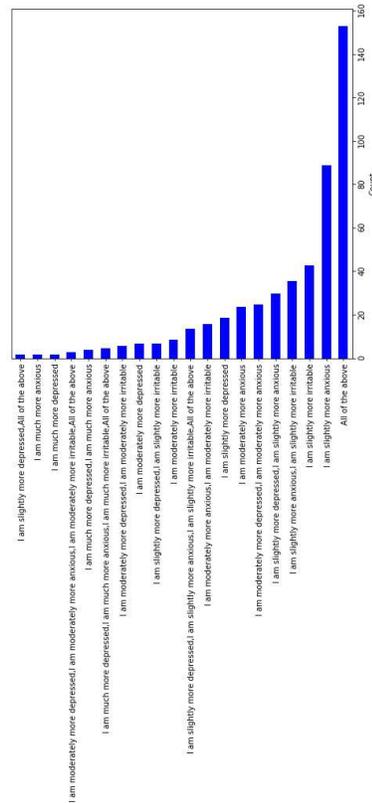
(C) Q27: Has the quality of your diet changed?



(D) Q28: Has the amount of food you have been eating per day changed?



(E) Q29 (a): Please tell us how your mood has changed. My mood has been:



(F) Q29 (b): Please tell us how your mood has worsened

FIGURE A.20 Frequency bar charts after the preprocessing phase