Illinois State University

ISU ReD: Research and eData

Theses and Dissertations

7-19-2021

COVID-19 Mortality Prediction Using Machine Learning-Integrated Random Forest Algorithm Under Varying Patient Frailty

Erwin Cornelius

Illinois State University, corneliuserwin84@gmail.com

Follow this and additional works at: <https://ir.library.illinoisstate.edu/etd>

Recommended Citation

Cornelius, Erwin, "COVID-19 Mortality Prediction Using Machine Learning-Integrated Random Forest Algorithm Under Varying Patient Frailty" (2021). *Theses and Dissertations*. 1485.

<https://ir.library.illinoisstate.edu/etd/1485>

This Thesis is brought to you for free and open access by ISU ReD: Research and eData. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of ISU ReD: Research and eData. For more information, please contact ISUReD@ilstu.edu.

COVID-19 MORTALITY PREDICTION USING MACHINE LEARNING-INTEGRATED RANDOM FOREST ALGORITHM UNDER VARYING PATIENT FRAILTY

ERWIN CORNELIUS

40 Pages

The abundance of type and quantity of available data in the healthcare field has led many to utilize machine learning approaches to keep up with this influx of data. Data pertaining to COVID-19 is an area of recent interest. The widespread influence of the virus across the United States creates an obvious need to identify groups of individuals that are at an increased risk of mortality from the virus. We propose a so-called clustered random forest approach to predict COVID-19 patient mortality. We use this approach to examine the hidden heterogeneity of patient frailty by examining demographic information for COVID-19 patients. We find that our clustered random forest approach attains predictive performance comparable to other published methods. We also find that follow-up analysis with neural network modeling and k-means clustering provide insight into the type and magnitude of mortality risks associated with COVID-19.

KEYWORDS: random forest; machine learning; COVID-19

COVID-19 MORTALITY PREDICTION USING MACHINE LEARNING-INTEGRATED
RANDOM FOREST ALGORITHM UNDER VARYING PATIENT FRAILTY

ERWIN CORNELIUS

A Thesis Submitted in Partial
Fulfillment of the Requirements
for the Degree of

MASTER OF SCIENCE

Department of Mathematics

ILLINOIS STATE UNIVERSITY

2021

Copyright 2021 Erwin Cornelius

COVID-19 MORTALITY PREDICTION USING MACHINE LEARNING-INTEGRATED
RANDOM FOREST ALGORITHM UNDER VARYING PATIENT FRAILTY

ERWIN CORNELIUS

COMMITTEE MEMBERS:

Olcay Akman, Chair

Dan Hrozencik

Bill Perry

ACKNOWLEDGMENTS

I would like to thank my advisor, committee members, and everyone else who helped make my work here possible.

E.C.

CONTENTS

	Page
ACKNOWLEDGMENTS	i
TABLES	iii
FIGURES	v
CHAPTER I: INTRODUCTION	1
CHAPTER II: MATERIALS AND METHODS	6
CHAPTER III: RESULTS	13
CHAPTER IV: DISCUSSION	19
CHAPTER V: CONCLUSION	32
WORKS CITED	35
APPENDIX A: CODE AVAILABILITY	40

TABLES

Table	Page
2.1. Features selected for analysis and their descriptions	6
3.1. Accuracy, Precision, Recall, and F1 Score for the prediction of patient mortality on the data used by Iwendi et al. [23]	13
3.2. Accuracy, Precision, Recall, and F1 Score for the prediction of patient mortality on CDC case surveillance data	13
3.3. Percentage of individuals who fall into health categories for overall sample of 10,000 and for each of the 4 clusters. Given as percentages. For the first three features, the number is not a percent, but the mean value	14
3.4. Percentage of individuals who fall into racial categories for overall sample of 10,000 and for each of the 4 clusters. AI/AN is American Indian/Alaska Native, and NH/PI is Native Hawaiian/other Pacific Islander. Given as percentages	14
3.5. Percentage of individuals who fall into age categories for overall sample of 10,000 and for each of the 4 clusters. Given as percentages	14
3.6. Percentage of individuals who fall into sex and geographic categories for overall sample of 10,000 and for each of the 4 clusters. Given as percentages	14
3.7. Neural network-estimated probabilities of death for a white male with no pre-existing conditions by region and age group. Given as percentages	15
3.8. Neural network-estimated probabilities of death for a black male with no pre-existing conditions by region and age group. Given as percentages	15
3.9. Neural network-estimated probabilities of death for a hispanic male with no pre-existing conditions by region and age group. Given as percentages	15

3.10. Neural network-estimated probabilities of death for an asian male with no pre-existing conditions by region and age group. Given as percentages	16
4.1. Average added risk of mortality from COVID-19 for membership in a certain feature level, as compared to the overall mortality rate of 5.62%. Given as percentages	25

FIGURES

Figure	Page
2.1. A decision tree classifier for classifying observations into three classes: 1, 2, and 3	8
2.2. An example of a neural network model	11
3.1. Density plot of observations projected onto the 2-dimensional plane. Top in red are deaths, and bottom in blue are survivors	17
3.2. Density plot of observations projected onto the 2-dimensional plane. Cluster 4 is the high death rate cluster	18
4.1. Percentage of individuals who fall into health categories for overall sample of 10,000 and for each of the 4 clusters. Same information as in last three columns of Table 3.3	21
4.2. Percentage of individuals who fall into racial categories for overall sample of 10,000 and for each of the 4 clusters. AI/AN is American Indian/Alaska Native, and NH/PI is Native Hawaiian/other Pacific Islander. Same information as in Table 3.4	22
4.3. Percentage of individuals who fall into age categories for overall sample of 10,000 and for each of the 4 clusters. Same information as in Table 3.5	23
4.4. Percentage of individuals who fall into sex and geographic categories for overall sample of 10,000 and for each of the 4 clusters. Same information as in Table 3.6	24
4.5. Probability of mortality for a white male with no pre-existing conditions for certain age groups in years. Given as a decimal. Same information as in Table 3.7	26
4.6. Probability of mortality for a black male with no pre-existing conditions for certain age groups in years. Given as a decimal. Same information as in Table 3.8	27
4.7. Probability of mortality for a hispanic male with no pre-existing conditions for certain age groups in years. Given as a decimal. Same information as in Table 3.9	28

4.8. Probability of mortality for an Asian male with no pre-existing conditions for certain age groups in years. Given as a decimal. Same information as in Table 3.10	29
4.9. Proportion of individuals in each age group by cluster. Cluster 4 was the high-mortality cluster	30
4.10. Proportion of individuals in each sex group by cluster. Cluster 4 was the high-mortality cluster	30
4.11. Proportion of individuals in each race group by cluster. Cluster 4 was the high-mortality cluster	31
4.12. Proportion of individuals in each medical condition group by cluster. Cluster 4 was the high-mortality cluster	31

CHAPTER I: INTRODUCTION

Healthcare in the United States is a constantly and rapidly-evolving industry. The extensive branches of the industry make the collection and integration of data an important endeavor. From development and testing of medical treatments or pharmaceuticals to insurance risk assessment and premium pricing, a wide variety of data is necessary for use by industry professionals.

A general trend in the current world of information is an increasing abundance and availability of data. This creates a need for the implementation of methods that can better process this abundance of data and turn it into useful information. This ever-evolving challenge requires continual development and integration of data analytics tools. To meet the needs of data analytics in the healthcare industry, machine learning (ML) approaches have become a topic of interest [1] [2] [3] [4] [5] [6] [7].

Unquestionably, the most significant challenge to face the healthcare industry in the United States over the past year has been COVID-19. Since the first confirmed case on January 21, 2020, the United States has seen over 32 million COVID-19 cases that have led to over 578,000 deaths [8]. Predicting the spread and severity of COVID-19 is an obvious area of interest.

A recent area of interest in the study of COVID-19 has been on those individuals at particular risk from COVID-19. According to Aizenman [9], treatment of individuals with compromised immune systems may shape the way the virus runs its course:

“There's mounting research to suggest that protecting people who are immuno-compromised from getting COVID is important not just for their sake – it could be critical in the effort to end the pandemic for everyone" [9].

For example, immuno-compromised individuals may not react to treatments in the same manner. In examining HIV-positive individuals who had been vaccinated against COVID-19, Dr. Laura McCoy, an infectious disease researcher at University College in London, remarked: "we couldn't actually see any measurable levels of anti-coronavirus antibody in the blood" [9]. These immuno-compromised individuals may also create avenues of risk for the general population. Dr. Salim Abdool Karim of South Africa's Centre for the AIDS Programme of Research noted that the increased time immuno-compromised individuals may be infected with COVID-19 creates added risk, noting that a particular HIV-positive woman "became a cauldron for the creation of a whole lot of new variants" [9]. It is apparent that identifying certain groups at an increased risk from COVID-19 could be important in curtailing the spread of the virus.

To aid in evaluating and modeling the course of COVID-19, ML methods have been employed. Wang and Wong [10] utilized a neural network approach on chest X-ray images to identify COVID-19 cases. Pal et al. [11] developed country-specific neural network models to determine a country's COVID-19 risk. In Liu et al. [12], the authors use ML approaches to predict the spread of COVID-19 in Chinese provinces. The authors in Beck et al. [13] propose a list of drugs to possibly combat COVID-19 identified from a deep learning model. A generative adversarial network was built by Khalifa et al. [14] to identify COVID-19 linked pneumonia from chest X-ray images. Sujath et al. [15] attempt to forecast the spread of COVID-19 in India using multilayer perceptron and vector autoregression methods.

The random forest (RF) is an approach in ML that has shown merit in accurately modeling COVID-19 outcomes. The RF approach utilizes an ensemble of weak decision tree classifiers to make predictions. The RF approach has been used somewhat commonly to model the COVID-19 pandemic. Tang et al. [16] built a RF model to predict the severity of a COVID-

19 diagnosis from chest CT images. Barbosa et al. [17] utilized a RF classifier to diagnose COVID-19 from blood samples. Gupta et al. [18] predicted cases of COVID-19 in India with a RF model. The RF model was also employed by Yesilkanat [19] to predict COVID-19 case counts at the country level. An effort to predict COVID-19 health outcomes from demographic and medical information in Korea was made by the authors in An et al. [20]. A similar effort to predict COVID-19 health outcomes for patients in Wuhan using RF methods was made by Wang et al. [21]. Majhi et al. [22] also utilized a RF model to predict COVID-19 case counts. Iwendi et al. [23] set out with the goal of identifying a modeling approach that worked best for COVID-19 patient health outcome prediction. The authors utilized the RF method with the AdaBoost boosting algorithm to predict COVID-19 mortality at the individual level. Demographic information for individuals that contracted the virus was used to predict if a patient would die. The authors compared several common modeling approaches and determined that the boosted RF model was the best.

While the RF method is generally praiseworthy for its prediction performance, it does have limits in explaining why certain predictions were made. The RF method does not assume a structure for the data, so it can be difficult to assess how and why certain features of the data affect predictive performance.

One challenge with COVID-19 is the somewhat varied effects that it can have on an individual. Whether it be individuals in different age groups or races, it is apparent that COVID-19 affects certain groups quite differently [24]. To this end, it could be beneficial to explore the differing risks that individuals have in the context of exposure to COVID-19. In the aforementioned work on COVID-19 patient health prediction, the models were developed using the full body of selected data. This could be in part due to the lack of available data. Given the

differing effects of COVID-19 on certain demographic groups, an exploration into modeling approaches that cater to these different effects seems worthy of investigation.

This idea of using individual patient characteristics to predict health outcomes underscores a current research trend in the healthcare industry: personalized medicine [25] [26] [27] [28]. The interest here is to use information about an individual to predict health outcomes.

To better predict a patient's health outcome, the main hook with an individualized model is the discernment with which a predictive model is built. Rather than using a predictive model that has been built using all available patients, a personalized model would be built using only those patients who are similar to the individual that is being assessed.

Building a model using only individuals similar to a patient in question obviously necessitates a way to define patient similarity. As the success of a personalized model may be heavily dependent on selecting appropriate individuals from which to build the model, appropriately defining a similarity measure is an important step.

Some different approaches have been utilized to measure patient similarity. Park et al. [29] used a so-called statistical case-based reasoning approach to identify similar patients. A clustering approach was used to identify similar patients by Panahiazar et al. [30]. A propensity score approach was utilized by Brookhart et al. [31] to define patient similarity.

Interestingly, a RF methodology can be employed to define a similarity measure. Lee [32] used RF methods to develop a patient similarity measure for ICU patients. Running a RF model in an unsupervised manner (with no prediction goal in mind) gives a similarity score between individuals. This similarity measure allowed a RF model to be developed for individual patients using only similar patients in the model-building process. The author found that the RF methods employed performed better than other common models to predict ICU patient mortality.

This paper intends to explain and implement a method to predict COVID-19 patient health outcomes using RF methods while concurrently giving insight into the different risks that different demographic groups face from COVID-19. As a note, we are eschewing more traditional approaches, such as principal components analysis, in favor of RF methods for the ability to better interpret our features. We will use the boosted RF model shown by Iwendi et al. [23] to be a reasonable standard for addressing prediction performance for COVID-19 patient mortality as a benchmark for predictive performance. We will also incorporate the RF patient-similarity measure utilized by Lee [32] as a tool to examine the qualities that affect COVID-19 risk. A neural network model and k-means clustering will both be used as methods for follow-up work to better understand the conclusions drawn from the RF modeling. This work will show that a RF model that utilizes preliminary clustering by patient similarity can achieve prediction performance competitive with other methods while also providing qualitative information about feature effects that are typically hard to get at with RF methods.

CHAPTER II: MATERIALS AND METHODS

Modeling for this project was done in R version 3.6.1. The *randomForest* [33] R package was used for building RF models. The *adabag* [34] and *cluster* [35] packages were used for implementing the AdaBoost algorithm and cluster analysis, respectively. The *neuralnet* [36] package was used for building the neural network model.

For this project, the data were sourced from the CDC's case surveillance data [37]. This dataset contains observations of over 22 million occurrences of COVID-19 in the United States with 32 features. The observations range in time from as early as January 1, 2020 to as late as March 30, 2021. A summary of these data tells us that only about 5% of individuals observed died from COVID-19. This unbalanced distribution of surviving and dying observations is important to note for putting future evaluation metrics into context. To ease the computational burden for this project, a sample of 10,000 individuals was randomly selected from the CDC data to be used in this project. Additionally, due to the presence of many missing values for some data features, a selection of 10 features were used for our analysis. The utilized features can be seen in **Table 2.1**. Additionally, date features were converted to numeric features to accommodate formatting for the utilized R packages.

Table 2.1: Features selected for analysis and their descriptions.

Feature	Description	Levels	Type
race_ethnicity_combined	Race and Ethnicity (combined)	8 levels	Categorical
cdc_case_earliest_dt	Earliest available date for record	NA	Date
cdc_report_dt	Initial date case reported to CDC	NA	Date
sex	Patient sex	Male, Female	Categorical
onset_date	Date of symptom onset	NA	Date
hosp_yn	Patient hospitalized	Yes, No	Categorical
death_yn	Patient died	Yes, No	Categorical
medcond_yn	Patient had pre-existing condition	Yes, No	Categorical

Feature	Description	Levels	Type
res_state	State of residence	50 levels	Categorical
age_group	Patient age group	8 levels	Categorical

Our method for predicting COVID-19 patient mortality in this project relies heavily on the random forest (RF) classifier from Breiman [38]. The RF classifier is itself made up of many decision trees. A decision tree classifier is made by successively splitting our data at decision nodes according to feature values. Our initial decision node splits the data into two groups according to a cutoff value for one of the data features. Then these groups are again split by a decision node, and this process continues, building out the “branches” of the decisions tree. When the splitting stops, the last remaining groups, or “leaves” of the decision tree provide the designation for which class individuals in that group belong. The feature used at each decision node to split the data is typically chosen so that error at that step is minimized. An example of a classification decision tree may be seen in **Figure 2.1**. The RF classifier uses an ensemble “forest” of these decision trees to make its classifications.

Figure 2.1: A decision tree classifier for classifying observations into three classes: 1, 2, and 3.

A decision tree classifier for classifying observations into three classes: 1, 2, and 3.

Each tree in the RF ensemble is built using a bootstrapped random sample of the available data and considering only a random selection of available features when each splitting node is made.

To classify an observation with the RF model, each decision tree in the ensemble “votes” for the class it predicts and the majority vote of the decision trees in the ensemble is the class that the RF classifier predicts. This reliance on the majority vote to classify an observation provides for better performance than a single decision tree classifier.

The authors in Iwendi [23] et al. elect to use boosting to improve the RF prediction results. The AdaBoost method the authors use works by building a sequence of RF models with misclassified observations given more weight for each RF model in the sequence. Overall classifications are given with an aggregate of the model iterations.

We take a different approach to improve the RF model. As Lee [32] did, the RF classifier can create a similarity measure if used in an unsupervised manner. This is done by building fake observations from the available data. The RF classifier is then used to classify the data as real or false. A similarity measure between any two observations from the real dataset is taken to be the proportion of the time those two observations ended up in the same leaf of a decision tree in the RF ensemble. From Lee, we may then denote the similarity between a patient i and a patient j , call it $SIM(i,j)$, by

$$SIM(i,j) = \frac{Prox_{i,j}}{\sum_{k=1}^N Prox_{i,k}}, i = 1, \dots, N, j = 1, \dots, N, i \neq j, i \neq k.$$

Here $Prox(i,j)$ is the number of trees in the RF ensemble that have both i and j in the same terminal node, or "leaf." If calculated for every pair of patients, we can develop a similarity measure for all considered observations.

This similarity measure can be used to cluster the data into similar groups. We will be using the PAM algorithm implemented with the *clara* function in the *cluster* R package to group our data. The appropriate number of clusters can be tuned by looking at the average silhouette for observations in a cluster. For each observation, the silhouette width is defined as the minimum average dissimilarity of that observation compared to each other cluster. Thus a high average silhouette width indicates that a cluster is different from the other clusters. That is, a high average silhouette width indicates that clusters are distinct, as desired.

We will now discuss the methodology that will be used for the analysis in this paper, which will be referred to as the "Clustered RF" method. The clustering described earlier will be used as a preliminary step to building our model. The clustering will be performed to partition the observations according to the highest average silhouette width, and then on each cluster, a RF model will be built to predict patient mortality. The classifications for the overall model will simply be the classifications given by the RF classifier built on each cluster. In this way, classifications for patients in a cluster will be based on information from the patients similar to them also in that cluster.

Given the skew of the data toward surviving observations, Iwendi [23] et al. chose to use several metrics to evaluate predictive performance. We will use those same metrics in our evaluation. These measures are based on True Positives (TP), True Negatives (TN), False Positives (FP), and False Negatives (FN) and are given now.

$$\begin{aligned}
 \text{Accuracy} &= \frac{TP + TN}{TP + TN + FP + FN} \\
 \text{Precision} &= \frac{TP}{TP + FP} \\
 \text{Recall} &= \frac{TP}{TP + FN} \\
 \text{F1 Score} &= 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}
 \end{aligned}$$

Given the small proportion of individuals who did die from COVID-19 in the available data, these measures give a better picture of how powerful a model is in correctly predicting mortality.

As we will see during our discussion of the results of the clustered RF method, the added risk that we are able to derive is a somewhat crude measure that does not take into account the risk associated with membership into multiple groups simultaneously. To get at the interaction between levels of different factors, a neural network model is used on the data as well. The

neural network approach is described in [39]. The neural network is built by deriving new features (these features are thought of as a hidden layer) that are linear combinations of the original features that have been passed through an activation function. Then the variable we wish to predict is modeled as a function of linear combinations of the derived features. A representation of the neural network model is seen in **Figure 2.2** (from [39]). The bottom row is the original data features.

Figure 2.2: An example of a neural network model.

The middle row is the hidden layer of derived features, and the top row is the variable we wish to predict. Fitting such a neural network model then requires the selection of the number of derived features in each hidden layer and the weights used to create the aforementioned linear combinations. Additionally, an appropriate activation function must be specified. The methodology for the fitting of the model is not a focus here, and we will not venture into the specifics. The takeaway here is that we will be using the neural network model to get the

probability that an individual with certain characteristics that we specify will die from COVID-19.

After observing the clusters that we obtain throughout the rest of this paper, a look at how discriminatory certain features of our data are in determining which individuals are clustered together will be a nice addition. To this end, we incorporate a k-means clustering of our data. This method clusters observations so as to minimize the sum-squared distance measure of each observation to its assigned cluster.

CHAPTER III: RESULTS

As a test of efficacy, the clustered RF method was performed on the data used by Iwendi [23] et al. using a partition of three clusters. The target feature to predict was patient mortality. To build each cluster's RF model, the *randomForest* command in the same-named R package was used with *ntree*=1000 and *mtry*=3. The evaluation metrics for the clustered RF method compared to the metrics obtained by the boosted RF method may be seen in **Table 3.1**.

Table 3.1: Accuracy, Precision, Recall, and F1 Score for the prediction of patient mortality on the data used by Iwendi et al. [23].

Model	Accuracy	Precision	Recall	F1-score
Boosted RF	0.94	1	0.75	0.86
Clustered RF	0.92	0.93	0.93	0.93

The clustered RF and the AdaBoost RF method were then performed on the CDC case surveillance data. Four clusters were used, as this number of clusters gave the highest average silhouette width of the considered values of 2:8. The target feature to predict was *death_yn*. To build each cluster's RF model, the *randomForest* command in the same-named R package was used with *ntree*=1000 and *mtry*=3. The evaluation metrics for these models may be seen in **Table 3.2**.

Table 3.2: Accuracy, Precision, Recall, and F1 Score for the prediction of patient mortality on CDC case surveillance data.

Model	Accuracy	Precision	Recall	F1-score
Boosted RF	0.95	0.59	0.40	0.48
Clustered RF	0.95	0.61	0.38	0.47

For the 4 clusters that our data was partitioned into, we have broken down the proportion of individuals that fall into each level of the features and compared them to the overall makeup of the sample of 10,000. We see these in **Tables 3.3, 3.4, 3.5, 3.6**.

Table 3.3: Percentage of individuals who fall into health categories for overall sample of 10,000 and for each of the 4 clusters. Given as percentages. For the first three features, the number is not a percent, but the mean value.

Cluster	cdc_case_earliest_dt	cdc_report_dt	onset_dt	pre-existing	hospital	died
Overall	266.04	281.77	267.58	46.34%	16.48%	5.62%
1	267.90	282.99	269.72	50.03	15.63	4.37
2	233.17	248.60	234.94	2.34	0	0
3	99.17	145.63	99.43	100	99.38	47.84
4	399.21	400.46	399.21	0	0	0

Table 3.4: Percentage of individuals who fall into racial categories for overall sample of 10,000 and for each of the 4 clusters. AI/AN is American Indian/Alaska Native, and NH/PI is Native Hawaiian/other Pacific Islander. Given as percentages.

Cluster	AI/AN	Asian	Black	Hispanic	Multiple	NH/PI	White
Overall	0.24%	2.88%	9.14%	20.4%	3.62%	0.42%	63.3%
1	0.24	2.74	9.45	21.46	3.52	0.46	62.13
2	0	0.78	1.17	8.20	2.73	0	87.11
3	0	4.94	24.69	23.46	4.01	0	42.90
4	0	2.26	3.39	14.25	2.94	0	77.15

Table 3.5: Percentage of individuals who fall into age categories for overall sample of 10,000 and for each of the 4 clusters. Given as percentages.

Cluster	0-9	10-19	20-29	30-39	40-49	50-59	60-69	70-79	80+
Overall	3.18%	9.8%	16.48%	15.08%	14.4%	15.2%	12.6%	8.08%	5.18%
1	2.92	7.61	14.95	15.24	15.72	16.72	13.02	8.63	5.18
2	2.15	28.71	46.48	10.74	6.25	3.71	1.76	0.20	0
3	0	0.31	3.09	4.94	6.48	21	26.54	15.74	21.91
4	11.54	30.32	18.55	22.62	12.67	4.07	0.23	0	0

Table 3.6: Percentage of individuals who fall into sex and geographic categories for overall sample of 10,000 and for each of the 4 clusters. Given as percentages.

Cluster	Female	Male	Northeast	Midwest	South	West
Overall	54.12%	45.88%	17.54%	53.86%	5.34%	23.26%
1	54.53	45.47	15.85	52.68	5.81	25.66
2	54.88	45.12	9.38	89.06	0	1.56
3	42.28	57.72	77.78	12.35	8.33	1.54
4	51.36	48.64	16.29	66.74	0	16.97

A neural network model was built for the CDC case surveillance data. Using the *neuralnet* command in R with *hidden=3*, *threshold=0.1*, and all others default. A logistic activation function was used with *death_yn* as the prediction target, coded 0 for no death and 1 for death. As a note, the goal of this follow-up was to compare mortality rates for a variety of different individual characteristics. During the initial run, the *hosp_yn* feature proved to be too strong of a factor in predicting death to make the results easily interpreted. Since we want to compare relative risk of individuals with different characteristics, this feature was removed for the neural network model results we will be discussing. Using the neural network model, a selection of mortality rates were calculated for individuals with characteristics of interest. These results can be seen in **Tables 3.7, 3.8, 3.9, 3.10**.

Table 3.7: Neural network-estimated probabilities of death for a white male with no pre-existing conditions by region and age group. Given as percentages.

Age	10-19	20-29	30-39	40-49	50-59	60-69	70-79	80+
Northeast	0.29	0.33	0.38	0.36	0.94	2.48	5.19	9.41
Midwest	0.31	0.33	0.41	0.34	0.93	2.04	3.26	5.36
South	0.02	0.36	0	0.56	3.62	21.09	8.51	28.46
West	0.32	0.33	0.48	0.33	1.21	2.49	3.05	4.32

Table 3.8: Neural network-estimated probabilities of death for a black male with no pre-existing conditions by region and age group. Given as percentages.

Age	10-19	20-29	30-39	40-49	50-59	60-69	70-79	80+
Northeast	0.32	0.33	0.7	0.35	2.35	5.11	6.18	8.7
Midwest	0.32	0.33	0.53	0.34	1.44	2.97	3.46	4.73
South	0.26	0.39	1.7	0.5	7.5	16.27	23.32	35.47
West	0.32	0.33	0.5	0.33	1.25	2.48	2.72	3.55

Table 3.9: Neural network-estimated probabilities of death for a hispanic male with no pre-existing conditions by region and age group. Given as percentages.

Age	10-19	20-29	30-39	40-49	50-59	60-69	70-79	80+
Northeast	0.3	0.34	0.59	0.36	1.92	4.44	6.59	10.44
Midwest	0.31	0.33	0.5	0.34	1.35	2.91	3.88	5.81
South	0.12	0.38	0.1	0.56	0.95	4.62	17.5	35.5

Age	10-19	20-29	30-39	40-49	50-59	60-69	70-79	80+
West	0.32	0.33	0.52	0.33	1.39	2.85	3.31	4.52

Table 3.10: Neural network-estimated probabilities of death for an asian male with no pre-existing conditions by region and age group. Given as percentages.

Age	10-19	20-29	30-39	40-49	50-59	60-69	70-79	80+
Northeast	0.29	0.34	0.67	0.38	2.51	6.14	9.07	14.68
Midwest	0.31	0.33	0.57	0.35	1.76	3.93	5.42	8.28
South	0	0.41	0	0.68	48.68	51.87	63.73	42.9
West	0.32	0.33	0.61	0.34	1.87	3.98	4.68	6.47

The *kmeans* R command was used to cluster our data into 4 clusters. The number of clusters was chosen both because it is the same number used in the RF clustering and because the useful distinction amongst clusters in regards to death rate seemingly began to deteriorate when more clusters were used. We project our data onto the two-dimensional axis, seen in **Figure 3.1**. The same data plotted with the four cluster labels is seen in **Figure 3.2**. Our interest in this clustering application was looking at which features of the data were most discriminatory in determining membership in a certain higher risk group. To this end, we plotted the proportion of certain feature levels by membership in the four clusters, which we see in **Tables 4.9, 4.10, 4.11, 4.12**.

Figure 3.1: Density plot of observations projected onto the 2-dimensional plane. Top in red are deaths, and bottom in blue are survivors.

Figure 3.2: Density plot of observations projected onto the 2-dimensional plane. Cluster 4 is the high death rate cluster.

CHAPTER IV: DISCUSSION

Our aim in the implementation of the clustered RF is not to sacrifice predictive power. Consequently, it is important that our clustered RF approach performs comparably to the AdaBoost RF method determined by Iwendi et al. [23] to be a benchmark for COVID patient predictive models. We see from the comparison metrics in Table 3.1 that the clustered RF method does seem to perform comparably to the boosted RF. The Boosted RF model is superior in Accuracy and Precision, but the Clustered RF performs better in Recall and F1-score. It is fair to say that the clustered RF method is not an obvious step backwards. This gives a positive indication that the preliminary clustering can replace the boosting algorithm and give similar prediction performance. The comparison metrics for the clustered RF and the boosted RF on the CDC data seen in Table 3.2 yield similar results. The clustered RF seems to perform comparably to the AdaBoost method, with all compared metrics being within 0.02 of each other.

We do note that the performance of both methods on the CDC data is obviously much worse than the performance on the data from Iwendi et al. [23]. This is a bit troubling, but a deeper dive into the data used by those authors provides a possible explanation. Upon inspection of the data used, it is noteworthy that the deaths in these data have strong correlation with the patient being from a certain location. For example, 39 of the 63 reported death in these data came from China, and those 39 deaths were out of the total of 42 observations from China. It is then possible that this and other similar trends in the data created an unrealistic picture of correlation between country of origin and mortality. This has possibly created the capacity for the RF method to provide a prediction accuracy beyond what might be expected from a larger dataset, such as that provided by the CDC.

Having established that the clustered RF method meets a reasonable standard of predictive performance, it becomes reasonable to now look at the benefit of the clustering approach. While RF methods typically provide a lesser ability to look at how certain predictors influence predictions compared to parametric methods, this clustering approach provides an avenue to gain some insights in this regard. The cluster analysis grouped our 10,000 observations into four groups. Naturally, some of these clusters will have mortality rates higher than the overall mortality rate for the 10,000 observations, and some of these clusters will have mortality rates lower than the overall. By examining how clusters vary by death rate, we can attempt to gain some information about the characteristics of individuals at a higher risk of death from COVID. These trends may be compared against trends that have been noted by the CDC to see if the insights given by the clustering RF model could be reasonably trusted.

Note that in **Table 3.3**, cluster 1 has a death rate of 4.37%, which is the closest of the four clusters to the overall death rate of 5.62%. This fits a general trend we see for cluster 1. For every feature in this project, cluster 1 stays close to the makeup of the overall sample. In no obvious way does cluster 1 distinguish itself from the whole. Consequently, we will be focusing on clusters 2-4 to gain insight into the variable risk of COVID-19. The mortality rates for individuals in Clusters 2-4 are drastically different from the overall mortality rate of 5.62%. Clusters 2 and 4 did not have any deaths, while cluster 3 had a death rate of 47.84%. The natural inclination then is to view clusters 2 and 4 as low-risk clusters, and cluster 3 as a high-risk cluster with the intention of discerning characteristics of the individuals who reside in these clusters.

Figure 4.1: Percentage of individuals who fall into health categories for overall sample of 10,000 and for each of the 4 clusters. Same information as in last three columns of **Table 3.3**.

First, let us turn our attention to the features from **Figure 4.1**. We note that clusters 2 and 4 both had no patients that were hospitalized and a very low percentage of patients with pre-existing conditions. Cluster 3 had close to 100% of individuals hospitalized and having pre-existing conditions. This is in line with the risk of pre-existing conditions outlined by the CDC [40].

Figure 4.2: Percentage of individuals who fall into racial categories for overall sample of 10,000 and for each of the 4 clusters. AI/AN is American Indian/Alaska Native, and NH/PI is Native Hawaiian/other Pacific Islander. Same information as in **Table 3.4**.

In **Figure 4.2**, we find the racial breakdowns for the four clusters. We can find a clear picture here. The high-risk cluster 3 is over-represented with individuals of Asian, Black, Hispanic, and multiple ethnicity and under-represented in White individuals. The low-risk clusters 2 and 4 show the opposite trend, both having more White people and fewer of all other groups. It is clear that minority groups are much more abundant in these high-risk clusters and obviously less abundant in the low-risk clusters. This is a compatible conclusion with some of the risk factors also outlined by the CDC [24].

Figure 4.3: Percentage of individuals who fall into age categories for overall sample of 10,000 and for each of the 4 clusters. Same information as in **Table 3.5**.

We can also look at the age demographics in **Figure 4.3**. A general trend to note is that clusters 2 and 4 have a lower percentage of individuals in age groups above 39 than the overall, while cluster 3 has more individuals in the age categories above 49. Cluster 3 also has fewer individuals from all age groups 49 and younger. The trends for clusters 2 and 4 are less obvious in the younger age groups, but the general trend of older individuals being at greater risk follows some of the notions put forth by the CDC [40].

Figure 4.4: Percentage of individuals who fall into sex and geographic categories for overall sample of 10,000 and for each of the 4 clusters. Same information as in **Table 3.6**.

Finally, in **Figure 4.4**, we have sex and geographic region breakdowns. Cluster 3 has relatively more males than overall and more people from the Northeast and South. There are also fewer people from the Midwest in cluster 3, while the percentage of people from the West is close to 0. Clusters 2 and 4 have sex proportions similar to the overall but noticeably have more people from the Midwest and no people from the South. Here we seem to have highlighted both being Male and being from the Northeast or South as possible added risk factors for COVID-19 mortality.

Given the absence of an assumed structure placed on our data by the RF method, calculating added risk factor by membership is not a straightforward task. However, we can obtain a somewhat crude measure by calculating how likely an individual is to find themselves in one of our clusters. For example, to calculate the added risk associated with being male, we look at how males are distributed among the clusters and calculate how likely a male is to fall into each cluster. We then multiply by the likelihood that an individual dies in each cluster to get how likely a male is to die. While this calculation ignores the interaction of other factors, it can give a

general idea of what is happening. We look at a calculation of added risk for a selection of factors in **Table 4.1**. We see similar trends to what we might expect with minority groups and older individuals having added risk, with males and the Northeast and South also being identified as risk factors to COVID-19 mortality.

Table 4.1: Average added risk of mortality from COVID-19 for membership in a certain feature level, as compared to the overall mortality rate of 5.62%. Given as percentages.

Factor Level	Added Risk
Asian	+1.11%
Black	+2.41%
Hispanic	+0.13%
White	-0.84%
Female	-0.56%
Male	+0.09%
Pre-existing Condition	+1.73%
50-59 yrs	+0.54%
60-69 yrs	+1.75%
70-79 yrs	+1.50%
80+ yrs	+4.65%
Northeast	+4.70%
Midwest	-1.54%
South	+0.95%
West	-1.31%

The method we used to estimate added risk from the RF clusters has the drawback of not providing an assessment of added risk for more than a single factor level at a time. For example, the RF calculation only provided an added risk of 4.7% for living in the northeast. It would be worth investigating if this risk affected all individuals from the northeast equally or if certain groups felt the detrimental effect of living in the northeast more strongly. To this end, we performed mortality prediction using the neural network model. Specifically, we looked at risk for individuals from several racial groups by region and age group. We looked at male

individuals with no pre-existing conditions to focus on the possible different effect that region would have on race and age group mortality. We have the mortality estimates for white, black, Hispanic, and Asian individuals in **Figures 4.5, 4.6, 4.7, and 4.8** respectively.

Figure 4.5: Probability of mortality for a white male with no pre-existing conditions for certain age groups in years. Given as a decimal. Same information as in **Table 3.7**.

One overarching trend is that the effect of region seems to be hidden in the earlier age groups. For all four races, the mortality rates for earlier age groups are indistinguishable by region. However, at the later age groups, we see that the mortality seems to increase for the south and northeast regions, compared to the midwest and west. We also note that for the younger age groups the higher risk for the black, Hispanic, and Asian groups that was detected earlier does not show up either. A reasonable explanation for this phenomenon would be that health outcomes for younger individuals are not dependent on medical intervention, and, consequently, any disparity in mortality due to differing quality or access to healthcare caused by membership in a racial or regional group does not show up for those younger individuals.

Figure 4.6: Probability of mortality for a black male with no pre-existing conditions for certain age groups in years. Given as a decimal. Same information as in **Table 3.8**.

Interestingly, we do see the influence of racial groups appear differently for the later age groups. The mortality rates for non-white individuals in the older age groups do tend to be higher than the mortality rates for white individuals in the same age groups. Additionally, the detriment associated with living in the northeast or south takes effect earlier in the non-white race groups. For example, a white individual does not see the northeast detriment until the age 70-79 age group, while the other racial groups see this departure start to occur in the 50-59 or 60-69 age groups. We also see that Hispanic individuals do not have especially worse mortality risk for living in the south until the 70-79 age group, contrary to the other age groups.

Figure 4.7: Probability of mortality for a hispanic male with no pre-existing conditions for certain age groups in years. Given as a decimal. Same information as in **Table 3.9**.

Another obvious trend is the unrealistically high mortality risk for the south for many of the older age groups. The south was one of the lesser reported regions in our data, and it is quite possible that we have some reporting bias, with possibly only the most severe cases being reported and contributing to these extreme predicted mortality rates. It should be kept in mind that the focus for these predicted mortality rates is to look at the heterogeneity of risk factors, rather than getting the most accurate mortality prediction.

Figure 4.8: Probability of mortality for an Asian male with no pre-existing conditions for certain age groups in years. Given as a decimal. Same information as in **Table 3.10**.

The general trend we see when plotting the data in a two-dimensional manner, as in **Figure 3.2**, is that the deaths are more heavily populated on the bottom left side of our plane. However, we do note that both deaths and survivors are found throughout the covered area of the plane with no real way to separate them into two neat clusters. Consequently, the features available to us are not discriminating factors in determining whether or not an individual dies from COVID-19. Despite this, clustering into 4 clusters, as seen in **Figure 3.3**, does get us somewhere. Cluster 4, seen toward the left of the plane in purple, is quite distinct from the other three clusters in terms of mortality. This cluster has a death rate of about 17%, while the other three clusters have death rates from around 2% to 4%. We can then look at how certain features are distributed amongst the 4 clusters to get a sense of what features determine membership in a cluster. We see these breakdowns in **Figures 4.9, 4.10, 4.11, 4.12**. The general trend we see is unsurprising. We see younger individuals under-allocated and older individuals over-allocated to the high-risk cluster. Also, males, non-white individuals, and individuals with pre-existing

conditions are over-allocated to the high-risk cluster. This makes sense, as our previous work has identified these as areas of risk for increased mortality.

Figure 4.9: Proportion of individuals in each age group by cluster. Cluster 4 was the high-mortality cluster.

Figure 4.10: Proportion of individuals in each sex group by cluster. Cluster 4 was the high-mortality cluster.

Figure 4.11: Proportion of individuals in each race group by cluster. Cluster 4 was the high-mortality cluster.

Figure 4.12: Proportion of individuals in each medical condition group by cluster. Cluster 4 was the high-mortality cluster.

CHAPTER V: CONCLUSION

In this project, we set out to predict COVID-19 patient mortality using demographic data from patients in the United States. The Random Forest method was an attractive machine learning approach that had already been utilized for similar purposes. In fact, the AdaBoost boosted random forest had been established as the method that gave superior predictive performance when compared to several common techniques. While predictive performance is important, we were interested in the heterogeneity of COVID-19 risk factors. That is, we wanted to look at risk factors beyond the typically noted ones, such as having a pre-existing medical condition or being in an older age group.

Given the lack of structure that RF methods impose on the data, getting this information required the incorporation of additional tools. To this end, we incorporated a patient similarity measure that was derived from using a RF classifier in an unsupervised manner. This gave us a similarity measure between patients that allowed us to cluster our observations into 4 clusters. By building a RF classifier separately on each cluster, we hoped to make up the predictive performance gained by using the AdaBoost algorithm while simultaneously figuring out a way to get at how certain factors put some groups at higher risk of death from COVID-19.

The clustering approach did seem to pay off in prediction performance, as our approach yielded prediction performance metrics that were all within 0.02 of the Adaboost method. Additionally, the clusters did provide insight into factors affecting risk of mortality from COVID-19. We were able to look at which clusters were at higher or lower risk of death and note which demographic groups were over or under-represented in these clusters. The identified factors did in general mesh with conclusions put out by the CDC.

To better get at the intricacies of the added risk measure produced by the clustered RF approach, we built a neural network model to predict patient mortality. When this model was used to estimate risk for individuals from different racial, regional, and age groups, we did expose some hidden heterogeneity. The risk associated with a certain region, for example, did not affect all races or age groups equally. This neural network model certainly exposed the complexity in estimating COVID-19 mortality risk.

With the wide range of outcomes associated with COVID-19 infection, a way to assign some kind of order to the severity of outcome was needed. To this end, we utilized k-means clustering, along with a way to project the multi-dimensional observations onto the two-dimensional plane. This approach did shed some light on how discriminatory our features were in predicting COVID-19 mortality. While our density plots did show a general trend of death observations showing up more heavily in one area, our observations were far from being able to be neatly clustered into a survivor group and a death group. However, the subsequent clustering did produce a distinctly high-risk cluster whose membership seemed to be determined by the identified mortality risk factors.

In addition to providing some insight into the hidden heterogeneity of COVID-19 mortality risk, this project has seemingly succeeded in showing that RF methods when approached from a different angle can provide valuable insight into how certain features of data affect the prediction problem at hand. Specifically, our results could be used as a guide for which communities most need resources to combat COVID-19. We see clear trends of older, minority individuals in the northeast and south as being at an elevated risk of mortality from COVID-19, and this information should prompt a focus on these individuals.

The ability to achieve good prediction performance and also get qualitative information about features of the data make a clustered RF approach worthy of exploration in similar problems. This approach could be used as a starting point to look at how certain features influence a dependent variable in cases where less is known at the outset. Further work in how the clustering is performed from the similarity measure developed from the RF method would be warranted. Only two methods to cluster the data were utilized in this project, and it is certainly possible that more creative ways to perform clustering might yield interesting results.

Disclaimer: The CDC does not take responsibility for the scientific validity or accuracy of methodology, results, statistical analyses, or conclusions presented.

WORKS CITED

- [1] A. M. Darcy, A. K. Louie and L. W. Roberts, "Machine learning and the Profession of Medicine.," *JAMA*, vol. 315, p. 551–552, February 2016.
- [2] D. W. Bates, S. Saria, L. Ohno-Machado, A. Shah and G. Escobar, "Big Data In Health Care: Using Analytics To Identify And Manage High-Risk And High-Cost Patients," *Health Affairs*, vol. 33, p. 1123–1131, July 2014.
- [3] V. Jain and J. M. Chatterjee, "Machine Learning with Health Care Perspective," 2020.
- [4] J. M. Chatterjee, "Bioinformatics using Machine Learning," *Global Journal of Internet Interventions and IT Fusion*, vol. 1, p. 28–35, 2018.
- [5] A. Khamparia, D. Gupta, V. H. C. de Albuquerque, A. K. Sangaiah and R. H. Jhaveri, "Internet of health things-driven deep learning system for detection and classification of cervical cells using transfer learning.," *The Journal of Supercomputing*, vol. 76, p. 8590–8608, 2020.
- [6] A. Waheed, M. Goyal, D. Gupta, A. Khanna, F. Al-Turjman and P. R. Pinheiro, "CovidGAN: Data Augmentation Using Auxiliary Classifier GAN for Improved Covid-19 Detection," *IEEE Access*, vol. 8, p. 91916–91923, 2020.
- [7] G. Sakarkar, S. Pillai, C. V. Rao, A. Peshkar and S. Malewar, "Comparative Study of Ambient Air Quality Prediction System Using Machine Learning to Predict Air Quality in Smart City," *Proceedings of International Conference on IoT Inclusive Life (ICIL 2019), NITTTR Chandigarh, India*, p. 175–182, 2020.
- [8] "CDC COVID Data Tracker," 2021.

- [9] N. Aizenman, "Protecting The Immuno-Compromised Against COVID Could Be Key To Ending The Pandemic," 2021.
- [10] L. Wang and A. Wong, "COVID-Net: A Tailored Deep Convolutional Neural Network Design for Detection of COVID-19 Cases from Chest X-Ray Images," *arXiv*, vol. 2003.09871, 2020.
- [11] R. Pal, A. A. Sekh, S. Kar and D. K. Prasad, "Neural Network Based Country Wise Risk Prediction of COVID-19.," *Applied Sciences*, vol. 10, August 2020.
- [12] D. Liu, L. Clemente, C. Poirier, X. Ding, M. Chinazzi, J. T. Davis, A. Vespignani and M. Santillana, "A machine learning methodology for real-time forecasting of the 2019-2020 COVID-19 outbreak using Internet searches, news alerts, and estimates from mechanistic models," *arXiv*, vol. 2004.04019, 2020.
- [13] B. R. Beck, B. Shin, Y. Choi, S. Park and K. Kang, "Predicting commercially available antiviral drugs that may act on the novel coronavirus (SARS-CoV-2) through a drug-target interaction deep learning model," *Computational and Structural Biotechnology Journal*, vol. 18, p. 784–790, March 2020.
- [14] N. E. M. Khalifa, M. H. N. Taha, A. E. Hassanien and S. Elghamrawy, "Detection of Coronavirus (COVID-19) Associated Pneumonia based on Generative Adversarial Networks and a Fine-Tuned Deep Transfer Learning Model using Chest X-ray Dataset," *arXiv*, vol. 2004.01184, 2020.
- [15] R. Sujath, J. M. Chatterjee and A. E. Hassanien, "A machine learning forecasting model for COVID-19 pandemic in India," *Stochastic Environmental Research and Risk Assessment*, vol. 34, p. 959–972, May 2020.

- [16] Z. Tang, W. Zhao, X. Xie, Z. Zhong, F. Shi, J. Liu and D. Shen, "Severity Assessment of Coronavirus Disease 2019 (COVID-19) Using Quantitative Features from Chest CT Images," *arXiv*, vol. 2003.11988, 2020.
- [17] V. A. de Freitas Barbosa, J. C. Gomes, M. A. de Santana, C. L. de Lima, R. B. Calado, C. R. B. Junior, J. E. de Almeida Albuquerque, R. G. de Souza, R. J. E. de Araujo, R. E. deSouza and W. P. dos Santos, "Covid-19 rapid test by combining a random forest based web system and blood tests," *medRxiv*, 2020.
- [18] V. K. Gupta, D. Kumar and A. Sardana, "Prediction of COVID-19 confirmed, death, and cured cases in India using random forest model," *Big Data Mining and Analytics*, vol. 4, p. 116–123, June 2021.
- [19] C. M. Yesilkanat, "Spatio-temporal estimation of the daily cases of COVID-19 in worldwide using random forest machine learning algorithm," *Chaos, Solitons & Fractals*, vol. 140, November 2020.
- [20] C. An, H. Lim, D. W. Kim, J. H. Chang, Y. J. Choi and S. W. Kim, "Machine learning prediction for mortality of patients diagnosed with COVID-19: a nationwide Korean cohort study," *Scientific Reports*, vol. 10, October 2020.
- [21] J. Wang, H. Yu, Q. Hua, S. Jing, Z. Liu, X. Peng, C. Cao and Y. Luo, "A descriptive study of random forest algorithm for predicting COVID-19 patients outcome," *PeerJ*, vol. 8, August 2020.
- [22] R. Majhi, R. Thangeda, R. P. Sugasi and N. Kumar, "Analysis and prediction of COVID-19 trajectory: A machine learning approach," *Journal of Public Affairs*, vol. e2537, November 2020.

- [23] C. Iwendi, A. K. Bashir, A. Peshkar, R. Sujatha, J. M. Chatterjee, S. Pasupuleti, R. Mishra, S. Pillai and O. Jo, "COVID-19 Patient Health Prediction Using Boosted Random Forest Algorithm," *Frontiers in Public Health*, vol. 8, July 2020.
- [24] "Risk for COVID-19 Infection, Hospitalization, and Death By Race/Ethnicity," 2021.
- [25] J. Andreu-Perez, C. Poon, R. Merrifield, S. Wong and G. Yang, "Big data for health," *IEEE J Biomed Heal informatics*, vol. 19, p. 1208, 2015.
- [26] M. J. Joyner and N. Paneth, "Seven Questions for Personalized Medicine," *JAMA*, vol. 314, pp. 999-1000, August 2015.
- [27] L. A. Celi, J. D. Marshall, Y. Lai and D. J. Stone, "Disrupting Electronic Health Records Systems: The Next Generation," *JMIR Med Inform*, vol. 3, p. e34, October 2015.
- [28] R. Xu, D. Nettleton and D. J. Nordman, "Case-Specific Random Forests," *Journal of Computational and Graphical Statistics*, vol. 25, p. 49–65, March 2016.
- [29] Y. J. Park, B. C. Kim and S. H. Chun, "New knowledge extraction technique using probability for case-based reasoning: application to medical diagnosis.," *Expert Systems*, vol. 23, p. 2–20, February 2006.
- [30] M. Panahiazar, V. Taslimitehrani, N. L. Pereira and J. Pathak, "Using EHRs for heart failure therapy recommendation using multidimensional patient similarity analytics.," *Stud Health Technol Inform*, vol. 210, p. 369–373, 2015.
- [31] M. A. Brookhart, S. Schneeweiss, K. J. Rothman, R. J. Glynn, J. Avorn and T. Sturmer, "Variable selection for propensity score models.," *Am J Epidemiol*, vol. 163, p. 1149–1156, June 2006.

- [32] J. Lee, "Patient-Specific Predictive Modeling Using Random Forests: An Observational Study for the Critically Ill," *JMIR Med Inform*, vol. 5, p. e3, 2017.
- [33] A. Liaw and M. Wiener, "Classification and Regression by randomForest," *R News*, vol. 2, pp. 18-22, 2002.
- [34] E. Alfaro, M. Gámez and N. García, "adabag: An R Package for Classification with Boosting and Bagging," *Journal of Statistical Software*, vol. 54, p. 1–35, 2013.
- [35] M. Maechler, P. Rousseeuw, A. Struyf, M. Hubert and K. Hornik, "cluster: Cluster Analysis Basics and Extensions," 2021.
- [36] S. Fritsch, F. Guenther and M. N. Wright, "neuralnet: Training of Neural Networks," 2019.
- [37] CDC, "Centers for disease control and prevention, covid-19 response. covid-19 casesurveillance data access, summary, and limitations (march 30, 2021 version)," 2021.
- [38] L. Breiman, "Random Forests," *Mach Learn*, vol. 45, p. 5–32, 2001.
- [39] T. Hastie, R. Tibshirani and J. Friedman, *The Elements of Statistical Learning*, p. 389–415, 2017.
- [40] "CDC People at Increased Risk," 2021.
- [41] S. Fritsch, F. Guenther and M. N. Wright, "neuralnet: Training of Neural Networks," 2019.

APPENDIX A: CODE AVAILABILITY

Code Availability: R code for this work may be accessed at
<https://github.com/edcorne/Cornelius-MS-Thesis-Summer-21-Illinois-State-University.git>.