

# Using Machine Learning to Discover Unexpected Patterns in Clinical Data: A Case Study in COVID-19 Sub-cohort Discovery

**Hannah Paris Cowley**

Johns Hopkins University Applied Physics Laboratory

**Michael S. Robinette**

Johns Hopkins University Applied Physics Laboratory

**Jordan K. Matelsky**

Johns Hopkins University Applied Physics Laboratory

**Daniel Xenos**

Johns Hopkins University Applied Physics Laboratory

**Aparajita Kashyap**

Johns Hopkins University Applied Physics Laboratory

**Nabeela F. Ibrahim**

Johns Hopkins University Applied Physics Laboratory

**Matthew L. Robinson**

The Johns Hopkins University School of Medicine

**Scott Zeger**

The Johns Hopkins University Bloomberg School of Public Health

**Brian T. Garibaldi**

The Johns Hopkins University School of Medicine

**William Gray-Roncal** (✉ [William.Gray.Roncal@jhupl.edu](mailto:William.Gray.Roncal@jhupl.edu))

Johns Hopkins University Applied Physics Laboratory

---

## Research Article

**Keywords:** Machine Learning , Discover Unexpected Patterns , COVID-19 Sub-cohort Discovery

**Posted Date:** January 6th, 2022

**DOI:** <https://doi.org/10.21203/rs.3.rs-1182055/v1>

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

[Read Full License](#)

---

# Using machine learning to discover unexpected patterns in clinical data: *a case study in COVID-19 sub-cohort discovery*

Hannah Paris Cowley<sup>1,\*</sup>, Michael S. Robinette<sup>1</sup>, Jordan K. Matelsky<sup>1</sup>, Daniel Xenos<sup>1</sup>, Aparajita Kashyap<sup>1</sup>, Nabeela F. Ibrahim<sup>1</sup>, Matthew L. Robinson<sup>2</sup>, Scott Zeger<sup>3</sup>, Brian T. Garibaldi<sup>3</sup>, and William Gray-Roncal<sup>1,\*</sup>

<sup>1</sup>The Johns Hopkins University Applied Physics Laboratory, Research and Exploratory Development Department, Laurel, MD, USA

<sup>2</sup>The Johns Hopkins University School of Medicine, Baltimore, MD, USA

<sup>3</sup>The Johns Hopkins University Bloomberg School of Public Health, Baltimore, MD, USA

\*Correspondence to: Hannah.Cowley@jhuapl.edu; William.Gray.Roncal@jhuapl.edu

## ABSTRACT

As clinicians are faced with a deluge of new information, data science can play a key role in highlighting key features towards developing new clinical hypotheses. Indeed, insights derived from machine learning can serve as a clinical support tool by connecting care providers with results from big data analysis to identify latent patterns that may not be easily detected by even skilled human observers. In this work, we show an example of collaboration between clinicians and data scientists during the COVID-19 pandemic, identifying sub-groups of COVID-19 patients with unanticipated outcomes or who are high-risk for severe disease or death. We apply a random forest classifier model to predict adverse patient outcomes early in the disease course, and we connect our classification results to unsupervised clustering of patient features that may underpin patient risk. The paradigm for using data science for hypothesis generation and clinical decision support, as well as our triage classification approach and unsupervised clustering methods to determine patient cohorts, are applicable to driving rapid hypothesis generation and iteration in a variety of clinical challenges, including future public health crises.

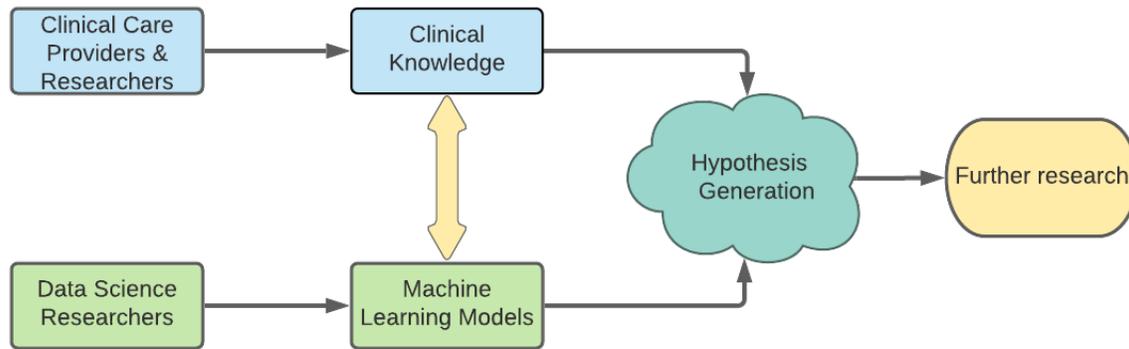
## Introduction

The integration of machine learning and data science analysis into clinical care represents a burgeoning opportunity to deliver even more impactful clinical care<sup>1-8</sup>, enabled in part by the vast amount of complex data that is collected as part of routine clinical encounters<sup>9,10</sup>. Inpatient hospital stays result in large amounts of patient data, including continuous monitoring of vital signs, frequent laboratory tests, clinical observations, radiology images, and more. These varied data modalities and their associated features are collected across a large patient population. While it might be challenging for even the most skilled clinician to synthesize and integrate the wealth of health data available into immediate clinical decision-making, these data are particularly valuable for assessing patterns in the population, such as identifying sub-groups of patients afflicted with a given disease. Leveraging big data to improve healthcare at an individualized level is the essence of precision medicine<sup>11</sup>.

Precision medicine and other big data modeling approaches can be used to augment clinical opinion, especially in the intensive care setting<sup>3,4,8,12,13</sup>. One of the most famous examples of this is the Acute Physiology and Chronic Health Evaluation (APACHE) model that predicts mortality of adult ICU patients using logistic regression with features derived from the first day of a patient's ICU stay<sup>14</sup>. Additionally, precision medicine tools can be used in a retrospective manner to find important features, covariates, and patterns in individual patient and patient sub-group trajectories that can be used prospectively in the care of future patients<sup>4,15-20</sup>.

Although precision medicine approaches are often tailored to a specific disease or use-case, the method we describe focuses on a readily generalizable model with a special focus on data exploration for identification of unexpected patterns and development of new avenues for scientific exploration. Specifically, we apply our model to COVID-19, but emphasize the potential utility of our approach in a variety of diseases — especially those for which timely diagnosis and treatment are critical for promoting patient health.

In this paper, we present a two-pronged approach to understanding COVID-19 disease presentation and progression. The first is a supervised machine learning algorithm, the “triaged” prediction model, designed to investigate features underlying patient outcomes and highlight patients who may have unexpected outcomes. A complementary, unsupervised machine learning



**Figure 1. Clinical knowledge and data science modeling work in tandem to drive hypothesis generation**

In our approach, we combine machine learning methods with clinical knowledge to generate hypothesized patient sub-groups and factors that affect patient outcomes. The collaboration between clinical insight and machine learning modeling informs the generation of new hypotheses which may then be tested in further research. The yellow arrow indicates the critical interplay between clinical knowledge and machine learning modeling: clinical knowledge informs the construction of such models and the results of models are checked against clinical knowledge.

clustering algorithm was built to find subsets of patients who present similarly at time of hospital admission and may represent clinically significant and distinct sub-groups. These analysis strategies utilize clinical input at all stages of research, from data collection to testing hypotheses generated from the machine learning methods.

The triaged prediction approach focuses on two goals: (1) identifying whether a patient will have a mild disease course or develop severe disease and/or death within fourteen days of hospital admission (where severe disease is defined as need for high-flow nasal cannula, non-invasive ventilation, or mechanical ventilation) and (2) identifying the patterns in patient variables responsible for differences in our ability to correctly classify patient outcomes. Unsupervised clustering is used to augment the prediction results by generating potential clinical sub-groups of patients and identifying which groups of patients may be more difficult to classify than others.

Through the methods we develop, we hope to provide clinicians with a data science framework to identify patients who are at an elevated risk of adverse outcomes from their disease and spur new thought into the patterns that might be causing this effect. A schematic of this approach can be found in Figure 1. Of particular importance is the philosophy of our approach: we use our models as hypothesis generation tools. Through analyzing the error patterns of imperfect models and using the modeling approach to identify patterns in patients with unexpected outcomes, we are able to pose potential avenues for further clinical investigation. These hypothesis generation tools are especially important in a pandemic setting, where diverse teams are working together to quickly find solutions for patients and when inequities in healthcare may be exacerbated. By building scalable, flexible data science tools relying on low-cost clinical resources, such as vital signs and simple blood laboratory tests, we can provide more equitable access to cutting-edge healthcare.

## Methods

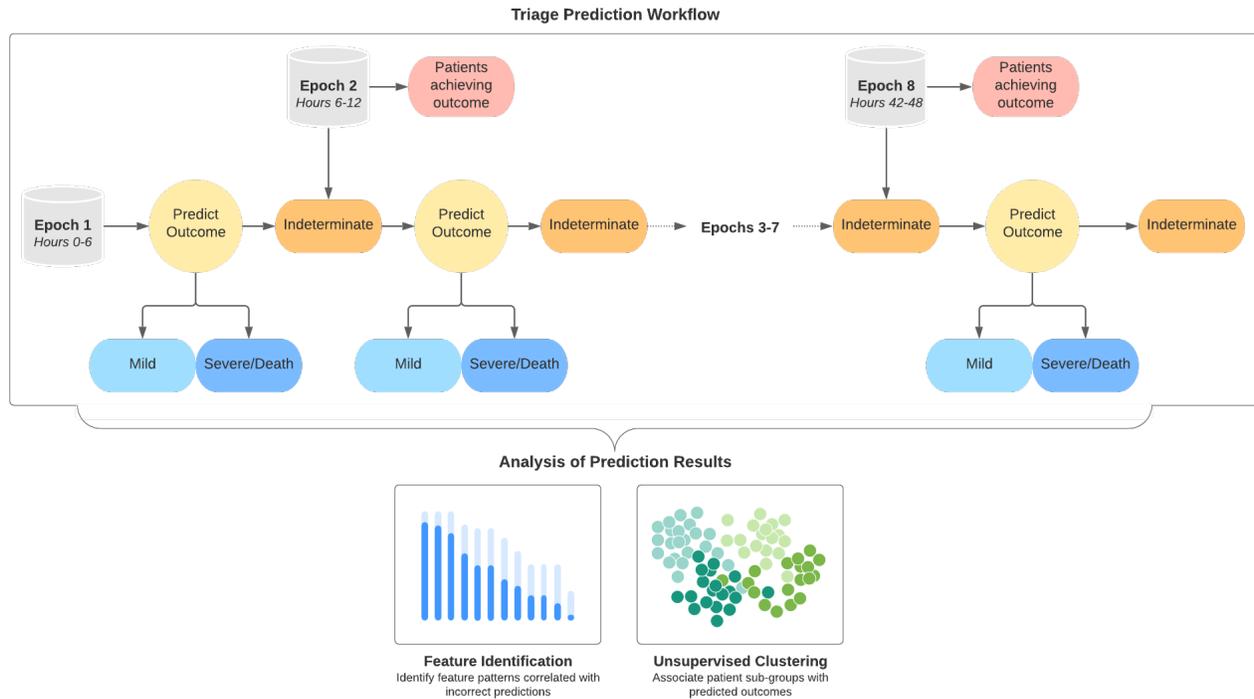
### Data

The data used were part of JH-CROWN: The COVID Precision Medicine Analytics Platform Registry<sup>21</sup> and some patients analyzed here have been included in previous descriptions of the cohort<sup>8,17-19,22-24</sup>. Data from 1175 patients in JH-CROWN were used in this analysis. This research was reviewed and approved by the Johns Hopkins Medicine Institutional Review Board (IRB) (IRB00250975), and all methods were conducted in accordance with the IRB guidelines and regulations for this style of research. Because this research posed minimal risk to participants and due to the retrospective nature of this work, the study was approved under a waiver of informed consent for participants.

Outcomes of interest in this study were defined as “mild” (hospital discharge without need for more intensive therapy than low-flow oxygen) or “severe/death”, where severe disease is defined as need for high-flow nasal cannula, non-invasive ventilation, or mechanical ventilation. For each patient, data from the first 48 hours after hospital admission were analyzed. The 48 hours were divided into 8 distinct 6-hour intervals, termed “epochs”. Patients were excluded from analysis if they achieved an outcome of interest within the first epoch, or if they did not achieve an outcome within 14 days of hospital admission. After

this down-selection, we were left with a cohort in which 25% of patients had a severe disease or death outcome within 14 days of hospital admission.

The data for each patient included demographics, presence or absence of comorbid conditions, COVID-19 symptoms present at time of admission, and vital signs and common blood laboratory measures per epoch. These features were selected because they are relatively easy to collect and accessible in a broad range of healthcare settings. A full description of the data has been published separately<sup>8</sup>.



**Figure 2. Triage Prediction Approach**

The goal of the triage prediction approach is to use the most minimal set of data possible to predict a given patient’s outcome and identify potential patient sub-groups based off of prediction patterns. At each epoch of analysis, patients are classified as either having a mild, severe disease and/or death, or indeterminate 14-day outcome. Patients who achieve an outcome during the epoch are removed from further analysis. After prediction, KL divergence is used to identify features that distinguish true positive from false negative predictions and the proportion of errors in each unsupervised cluster are calculated.

### Triaged Prediction Model

Before designing our own model, we examined the current research in generalizable modeling for COVID-19 patient prediction. Iwendi et al. (2020)<sup>25</sup> successfully applied a random forest classifier to the problem of identifying patients infected with COVID-19 based on symptoms in the outpatient setting. We built upon this approach for the epochs of our inpatient data and were informed by this approach in building our models.

The supervised model we designed was a triaged prediction approach, where we aimed to use the minimal time window of data possible to determine patient outcomes with high confidence. When an outcome could not be predicted with high confidence, additional data were added in 6-hour increments to augment the dataset available to the classifier. An illustration of the triaged prediction approach workflow can be found in Figure 2.

At each epoch, a random forest classifier was used to predict the patients’ 14-day outcome. Training and testing was accomplished using 10-fold cross-validation. The classifier could predict a mild outcome, a severe disease/death outcome, or an indeterminate outcome. If a patient had an indeterminate outcome, this meant that the data available for the patient up to and including the current epoch was not sufficient for confidently categorizing them as either severe disease/death or mild. For these patients, the next epoch of data was used and the prediction was re-attempted.

The thresholds for predicting a mild outcome or a severe disease/death outcome, defined as the proportion of decision trees

voting for the outcome of interest in the random forest, were 0.9 and 0.5 respectively. These thresholds were chosen after a sweep of positive class thresholds from 0.5 to 0.95 with the goals of reflecting the clinical need for high levels of certainty when declaring that a patient will have a mild outcome, and the desire to limit false negative predictions when determining that a patient will have a severe and/or death outcome after fourteen days.

Patients were removed from the epoch-by-epoch analysis pipeline after either (1) a confident prediction could be made about their fourteen-day outcome (regardless of the accuracy of this prediction) or (2) if the patient achieved an outcome of interest in the epoch (severe disease and/or death, or discharge from the hospital).

Although accurate prediction of patient outcomes is a main goal of the triaged prediction approach model, an equally important goal is to identify patients for whom these prediction methods fail - either by predicting an incorrect outcome or by being unable to render an outcome by the end of the 8 epochs of analysis. After predictions are made, we analyze the characteristics of patients for whom false negative and true positive predictions are made (incorrectly predicting a mild outcome and correctly predicting severe disease or death, respectively). The characteristics of these patients are then used, along with the results of unsupervised clustering of patients, as inspiration for new hypotheses about why particular patients have unpredictable outcomes.

### Unsupervised Clustering Model

Clustering analysis was used to identify patient sub-groups at time of hospital admission with the goal of investigating clinical features that may underpin different disease courses (Figure 3). The data for performing the unsupervised clustering approach included the first epoch of data as discussed above. Data were min-max normalized, placing features like age, vitals, and labs on the same 0-1 scale as the binary variables encoding comorbidities and symptoms. Dimensionality reduction to two dimensions was then performed via Uniform Manifold Approximation and Projection (UMAP) nonlinear embedding<sup>26</sup> in order to build a low-dimensional representation of our patient cohort. It is important to note that nonlinear embeddings like UMAP can cause issues with distance-based clustering and scoring methods, since densities of data points can be changed in the embedding process. We made attempts to mitigate this problem by enforcing a large number of neighbors in the UMAP embedding, which imposes more global structure in our reduced dataset<sup>26</sup>.

Clustering analysis was performed using hierarchical clustering with Ward's linkage<sup>27,28</sup> for a specified 20 clusters. 20 clusters was chosen based on two goals. First, we sought to identify a number of clusters that locally maximized the Calinski-Harabasz (CH) index<sup>29</sup>, the ratio of intercluster to intracluster dispersion. After sweeping over cluster numbers from 1 to 50, we identified that 20 clusters was a local maximum.

In addition to use of CH index, we favored a clustering representation that captured a range of disease severity, marked by the proportion of patients in the cluster with an outcome of severe disease or death. Having a broad range of clinical profiles and disease severity was thought to provide the greatest utility in the clinical setting, because then patients in each cluster might later be mapped to a prognosis and action plan.

After clustering, cluster labels were mapped to the original data for analysis of sub-group feature differences. We then analyzed the cluster feature distributions to find defining characteristics of the cluster. To find features that are most powerful for differentiating clusters, symmetric Kullback-Leibler (KL) divergence<sup>30</sup> calculations were performed on the feature distributions for each pair of identified clusters. The pairwise KL divergence measures were then summed to yield a total KL divergence for each feature. Features that had higher KL divergence information-gain values were then interpreted as being more important for separating clusters overall.

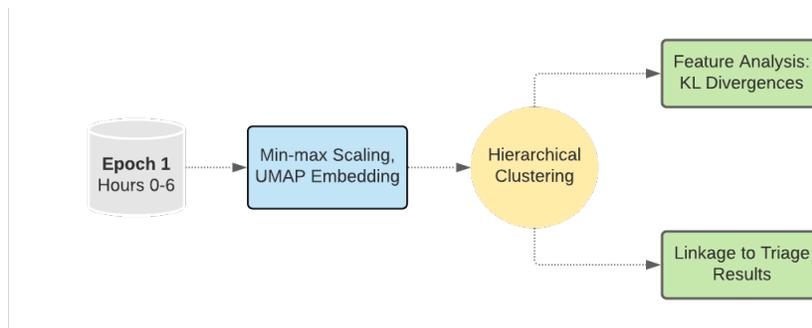
## Results

### Triaged Prediction

Of the 1175 patients in the cohort, 156 patients continued to have indeterminate predictions after eight epochs (48 hours) of analysis. A confusion matrix of prediction results by epoch of analysis with the number of true positive, false positive, true negative, and false negative predictions, and number of indeterminate and patients removed from analysis for achieving an outcome of interest before a prediction could be made, can be found in Table 1. Across all epochs, the false negative rate is low and the indeterminate rate is high - a product of our design choice to avoid making erroneous predictions.

The mean accuracy for prediction of 14-day outcome across the six-hour epochs (10-fold cross validation for each epoch) was 0.83 ( $\sigma = 0.08$ ), and the mean Matthew's correlation coefficient (MCC) was 0.24 ( $\sigma = 0.29$ ). Accuracy, AUC, and MCC per epoch of analysis can be found in Table 2. Of particular note is the high variability in MCC, a result of the predictions becoming more error-prone in later epochs of analysis. This is expected behavior due to the smaller numbers of patients eligible for analysis in later epochs. Additionally, later epochs of analysis may have more difficult patients to predict because the only patients present for prediction in these epochs are ones for whom a prediction could not be made previously.

The triaged classification system is imperfect at predicting 14-day outcome, but has possible utility for uncovering sub-groups of patients who have characteristics making them prone to mis-classification or indeterminate outcome. To investigate



### Figure 3. Unsupervised Clustering Approach

Clustering analysis was performed using the first epoch of triage analysis data. The goal of this complementary model was to find sub-groups of patients at the time of admission to the hospital. Data were preprocessed using min-max normalization and UMAP embedding to two dimensions. Hierarchical clustering using Ward's linkage was then applied to generate patient sub-groups. After clustering, the correlation of cluster membership with patient features and predictions from the triage prediction model were analyzed towards the goal of further identifying patient sub-groups and hypothesis generation.

the differences between those patients for whom a true positive outcome prediction (i.e., the model correctly predicted that the patient develops severe disease or death within the first 14 days of hospital admission) and false negative prediction (i.e., the model incorrectly predicted that a patient has a mild fourteen-day outcome) were made, we analyzed which features were most informative for separating out these patients in the first two epochs of analysis via KL divergence<sup>30</sup>.

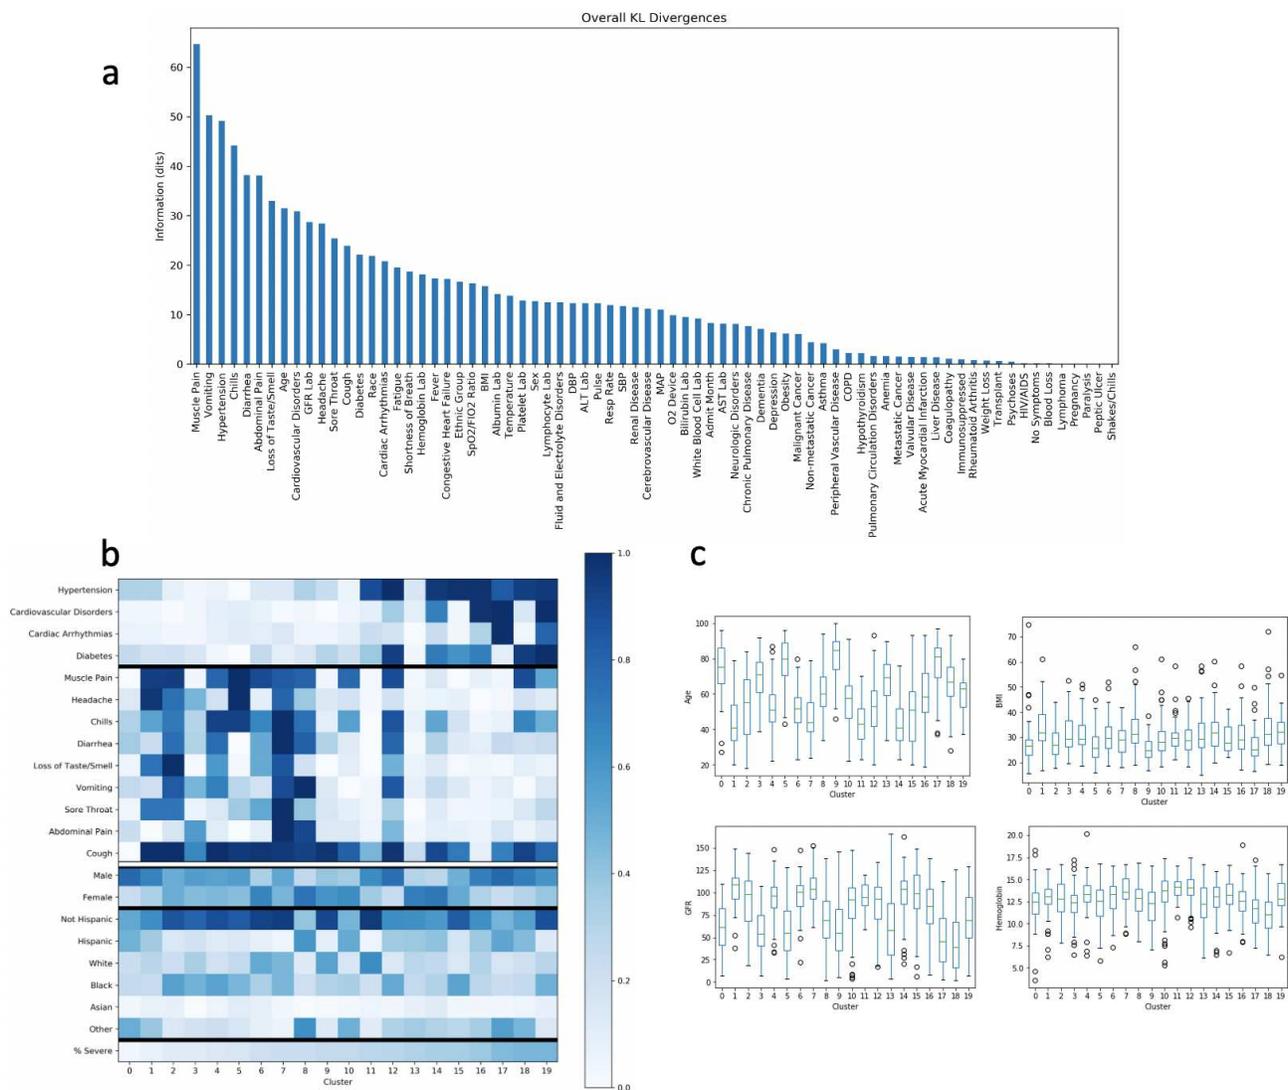
The first two epochs were the only epochs out of the eight with more than 5 patients having true positive and false negative prediction outcomes. Therefore, we analyzed these two epochs for the features that distinguish the true positive and false negative patient groups from each other. Through KL divergence calculations, we found that the most important features for distinguishing the false negative and true positive patient groups were SpO<sub>2</sub>/FiO<sub>2</sub> ratio, albumin and GFR lab test results, and the oxygen device that the patient used. In epoch 2, the same set of features in epoch 2, as well as the values for those features from epoch 1, were found to be important for distinguishing between the groups. The differences in feature values between the groups are illustrated in Table 3 and Table 4. Across epochs 1 and 2, patients for whom a false negative prediction was made appear to have significantly higher SpO<sub>2</sub>/FiO<sub>2</sub> ratios, as well as GFR and albumin test results that may fall into normal range during the first epoch of analysis.

Overall, the patients for whom a false negative prediction is made seem to have more mild symptoms during the first 2 epochs (12 hours) of analysis. The set of patients for whom a false negative error is made may represent a particularly interesting sub-group of patients for further investigation. Our unsupervised clustering is one way to investigate those features that are associated with an increased risk of erroneous or indeterminate outcomes.

### Unsupervised Clustering

Our clustering approach found patterns in patient sub-groups that may indicate potential clinically relevant differences in disease progression, with respect to a given patient's disease presentation at admission. By design, the 20 clusters identified exhibit a wide range of disease outcomes: each cluster is comprised of anywhere from 2% to 46% of patients achieving the severe disease/death outcome, as seen in Figure 4b. Via KL divergence, specific features were identified that drive separability of the 20 clusters. These features fall into the following categories: presence or absence of specific admission symptoms and comorbidities, age, race, ethnic group, and lab test results. Admission symptoms found to be important for separating clusters included muscle pain, headache, chills, diarrhea, loss of taste/smell, vomiting, sore throat, abdominal pain, and cough. Interestingly, this list includes all available symptoms except for fever, fatigue, shortness of breath, and "no symptoms" (kept separate to differentiate whether a patient affirmed that they had no COVID-19 symptoms, versus whether no symptoms were entered by the provider). Specific comorbidities important for separating clusters included hypertension, cardiovascular disease, cardiac arrhythmias, and diabetes. For visualization purposes, we separated out binary variables, which are expressed as proportion of individuals having the feature within the cluster, from continuous variables, expressed as the distribution of the feature within the cluster. Continuous features found to be important included age, BMI, GFR, and hemoglobin lab tests. A visualization of feature KL divergences can be found in Figure 4(a).

Clusters with a higher proportion of severe outcomes tended to have a higher proportion of comorbidities and lower proportion of reported symptoms when compared to clusters with a higher proportion of mild outcomes. However, our analysis finds more fine-grained sub-groups of patients for which this general trend is not the case. For example, cluster 13 has a high



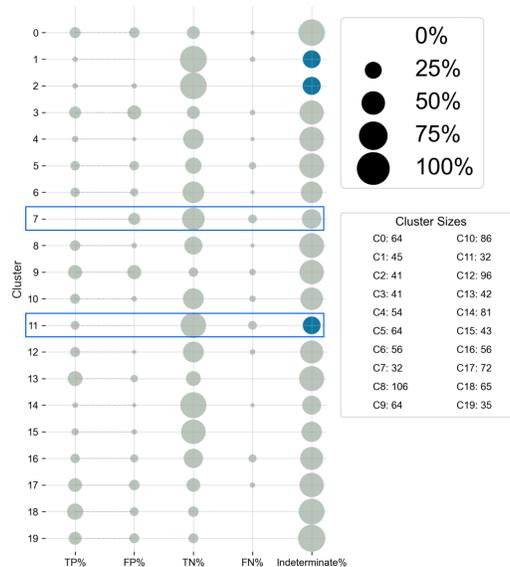
**Figure 4. Feature differences among clusters** a) Cluster-wise symmetric KL Divergences, sorted by information gain. Taller bars indicate that the feature is more important for distinguishing between unsupervised clusters. b) Heatmaps from top 20 KL divergences for binary features (comorbidities, symptoms), sorted by KL divergence values. Plotted values are proportions of present features for a given cluster, ranging from 0 to 100%. Demographics are also plotted between clusters for reference. Colors represent proportions of cluster makeup for each demographic. Clusters are sorted by the prevalence of severe disease or death outcomes within the cluster (bottom row). c) Boxplots for age, BMI, GFR, and hemoglobin. These data are presented as boxplots instead of in the heatmap fashion in subplot b) given the continuous nature of the data.

incidence (33%) of severe outcomes but low proportions of both comorbidities and symptoms. This group of patients does, however, tend to be higher in age than other patient clusters, which could be a driving force for worse outcomes. Clusters 7 and 11 also have higher and lower incidences of symptoms, respectively, when compared to neighboring clusters. Patient age, previously identified by researchers to drive patient outcomes (e.g.<sup>31</sup>) does not appear here to have a clear relationship with cluster membership (and thus a relationship with proportion of individuals having a severe and/or death outcome) when taken in isolation. A visualization of categorical and continuous variables and their relationship to cluster membership and proportion of individuals achieving a severe and/or death outcome can be found in Figure 4(b and c).

Finally, we analyzed patterns between cluster membership and triage prediction results, examining the proportion of patients in each cluster for whom a true positive, true negative, false positive, false negative, or indeterminate prediction was made. The results of this analysis can be found in Figure 5. Clusters 1, 2, and 11 all have less than 30% of the cluster remaining with an indeterminate prediction after 48 hours, denoted by a blue colored circle in the figure. For patients in these clusters in

particular, the triaging approach as described seems particularly well-suited for determining patient outcomes with certainty. Note, however, that high certainty does not always map to accurate predictions, such as in cluster 11.

There are two clusters, indicated by blue boxes, for which greater than 5% of patients had a false negative classification result: clusters 7 and 11. Cross-referencing these clusters with the symptom information displayed in Figure 4 (b), we found that these sub-groups with high false negative classifications also were noted to be exceptions to the general trends seen in other sub-groups, where high incidence of symptoms is related to better outcome. In particular, cluster 7 has high incidence of all symptoms found to be important, while cluster 11 has very few patients with incidence of these symptoms. For these clusters of patients, additional features not used in our triaging approach are likely necessary to make accurate predictions about disease outcome. Patients hypothesized to belong to these clusters should be treated with extra care when predicting outcomes, as they may appear to be on track for a good outcome early in their disease course, but instead achieve a severe and/or death outcome.



**Figure 5. Triaged classification results by unsupervised cluster**

The cluster membership of patients was cross-referenced with the final prediction result obtained, connecting the patient presentation within the first epoch of analysis with the final prediction outcome. Sizes of circles represent the proportion of patients within the cluster with the specified classification results. Visualizing the data in this manner aids in finding sub-groups of patients who are particularly well-suited for the triaged classification approach and driving discussion about factors influencing classification results.

## Discussion

In this work, we have shown how two complementary data science models are used, together with clinical insight, to identify patterns in patient characteristics and clinical features that may lead to differences in COVID-19 disease presentation and outcomes. Through our triaged prediction model, we ground our search for patient sub-groups in predicting whether patients will have mild or severe and/or death outcomes. In our clustering model, we search for patterns in patients as they first presented to the hospital, with a particular focus on investigating how characteristics of these unsupervised groupings of individuals might contribute to the ease or difficulty with which an outcome can be predicted.

One important aspect of this work is our paradigm for combining clinical and data science knowledge together to build models that can drive hypothesis generation. In our specific implementation, we show the power of partnership between data scientists and clinical experts when working to create new precision medicine tools and generating new avenues for investigation. Our two-pronged triage outcome prediction and data-driven sub-population discovery procedure may be easily extended to other disease models and domains. In this way, data science can be employed as a tool that “points clinical researchers to the *right* haystack” for further exploration and analysis. This is an especially useful approach in the context of emerging disease threats where existing knowledge is limited.

The unsupervised clustering results indicate that clinically-relevant patient sub-groups may be found in COVID-19 using readily-available patient data. Importantly, these sub-groups are characterized by the complex interactions between variables in patient health records. These sub-groups have different characteristics at time of admission and different proportions of severe

and/or death outcomes. It may be more difficult to accurately predict outcomes for some sub-groups of patients, as shown via the combination of triaged prediction results and the unsupervised clustering analysis. Although the exact reasons for the differences in these subgroups are beyond the scope of the present study, analyses such as these may lead to the generation of clinical hypotheses that are valuable in public health emergencies. We note that our dataset was drawn from early admissions during the COVID-19 pandemic and due to the emergence of different disease variants, vaccines, and human behaviors, these results may not generalize to contemporary cases.

Our clustering analysis produced discussion across our collaborative clinical and data science team as we sought to explain why the clustering results indicated in general, that clusters with a higher percentage of severe patients had a lower incidence of symptoms like diarrhea, loss of taste or smell, vomiting, sore throat, and abdominal pain. One hypothesis is that these patients enter the hospital when they are so severely ill that these symptoms remain undocumented, either because a patient cannot express these symptoms themselves, or because a patient is quickly attended to before this information is added into the medical record. The disease profiles of the milder patients with high prevalence of symptoms may also indicate more active immune systems, which cause worse symptoms early in a patient's disease course but a less severe outcome overall. These hypotheses could be used to investigate clinical care practices or some underlying factor in the disease model not yet uncovered.

Some other areas for future investigation are highlighted by this work, such as why the patients with false negative predictions in the first two epochs seemed to have more positive clinical profiles during the first two epochs of analysis when compared to their counterparts with true positive predicted outcomes, or how the addition of other variables, such as socioeconomic status and patient location, impact sub-group identification. Of potential interest to a clinical care provider might be why some patients continue to have an indeterminate prediction even after 48 hours of analysis, and how to accurately treat these patients for the best odds of survival. While this work cannot find answers to these questions, it is this style of modeling that may aid clinical research teams in partnership with data science teams in identifying clinical patterns that were previously not observed.

Like all models, our specific implementation comes with limitations. We cannot infer a causal relationship between the features of interest that separate patient sub-groups, since our models infer statistical dependencies only. There are limitations to the dataset used as well; we did not have access to the time since symptom onset, which might impact results because some patients may present to the hospital at different times in their COVID-19 disease course for different reasons. Additionally, although we emphasize the generalizability of our overarching approach of pairing supervised machine learning with unsupervised clustering for hypothesis generation and data exploration, our specific model implementations may not perform well on all disease models. Choosing appropriate models and then tuning models to a different disease or sub-group identification problem will require both clinical and data science expertise. We note that this data was taken from a particular time period of the COVID-19 pandemic, which continues to rapidly evolve; therefore, specific conclusions would need to be validated on current patient data prior to clinical use.

In the current study, the data were carefully collected, quality controlled, and stored in an easily-accessible format. Access to this readily-available, high-quality data source surely contributed to our success in developing this work. Approaches like ours, which draw on collaborative data-collection efforts and standardization protocols, have great potential for improving treatment efficacy and enabling clinicians to make different decisions when allocating limited resources and clinical care. This is a powerful argument for data standardization and co-location of compute infrastructure with clinical data<sup>23,32</sup>, *before* the next public health crisis arises.

In the future, precision modeling will likely continue to provide ground-breaking predictive modeling and sub-population discovery insights. These approaches, additionally, may provide an interface between emerging data science and clinical judgement. In this way, the data science underpinning precision medicine may serve as a hypothesis-generation engine, which can speed up clinical progress, and ultimately provide better care to patients, especially in the face of rapidly changing clinical landscapes.

## Acknowledgements

The authors thank and acknowledge Fatima Koroma, Olayide Ashiru, and Emma Liu for their contributions to the framing of this work and their helpful editing of this manuscript. Thank you to JHU C-STAT team members Paul Nagy, John Muschelli, Joshua Betz, and Karen Bandeen-Roche, for their review of this work. Gratitude is additionally extended to Johns Hopkins University and Johns Hopkins University Applied Physics Laboratory PMAP team members, including Suma Subbarao and Alan Ravitz, for their support of the Precision Medicine Analytics Platform (PMAP) program. The data utilized for this publication were part of the JH-CROWN: The COVID PMAP Registry which is based on the contribution of many patients and clinicians.

## Author contributions statement

HPC led development of this manuscript and the modeling approaches described herein. MSR co-lead clustering model development and contributed to the manuscript text. JKM contributed to the development of software tools for data access. DX contributed to model development and baseline classification results. AK and NFI contributed to the manuscript text and analysis efforts. MLR lead curation of the data used in this study. BTG provided clinical feedback to the experimental results and analysis. SZ provided statistical feedback and overall framing for the research study, along with BTG. WGR supervised this work and provided technical feedback to experimental results and analysis. All authors reviewed and approved the final manuscript.

## Additional information

Dr. Garibaldi is a member of the FDA Pulmonary-Asthma Drug Advisory Committee and has received consulting fees from Janssen Research and Development, LLC, Gilead Sciences, Inc and Atea Pharmaceuticals, Inc. All other authors declare no competing interests.

## Funding

This study received funding from from Hopkins inHealth and the Johns Hopkins Precision Medicine Program.

## Tables

**Table 1. Prediction results for each epoch of analysis**

Guided by clinical input, the triaged prediction algorithm was designed to predict outcomes with a low false negative rate at the expense of a higher indeterminate prediction rate. Patients were removed from the prediction process if they achieved an outcome of interest. The trade-off between false negative and indeterminate prediction results in 156 patients for whom a prediction could not be made within 8 epochs of analysis. For different clinical use-cases, the thresholds for making a positive or negative prediction could be adjusted.

Epoch	TP	FP	TN	FN	Indeterminate	Removed For Achieving Outcome
1	75	36	243	7	814	0
2	16	15	55	7	622	99
3	14	6	32	2	495	73
4	2	2	18	5	413	55
5	0	1	16	1	330	65
6	2	3	12	1	234	78
7	0	2	10	2	195	25
8	2	0	0	2	156	25

**Table 2. Mean metrics for each fold of 10-fold cross-validation over the 8 6-hour epochs under study (48 hours total).** After 3 epochs, performance drops as marked by the stark decrease in the Matthew’s Correlation Coefficient Metric, illustrating the difficulty of predicting outcomes for the subset of individuals that have indeterminate predictions in the preceding epochs.

Epoch	Accuracy	AUC	MCC
1	0.88	0.90	0.71
2	0.78	0.81	0.39
3	0.85	0.91	0.58
4	0.77	0.45	0.17
5	0.90	0.50	0.00
6	0.83	0.75	0.15
7	0.67	0.00	-0.10
8	0.92	0.75	0.05

**Table 3. Key feature differences between TP and FN groups predicted during epoch 1**

SpO<sub>2</sub>/FiO<sub>2</sub>, albumin, estimated glomerular filtration rate (GFR) test results, and mode of oxygen delivery during epoch 1 were found to be the most important for distinguishing patients for whom a false negative and true positive prediction was made at this stage of analysis. Feature importance was calculated via KL divergence and statistical testing for differences of means for laboratory and vital signs were conducted via ANOVA (\* p<.05, \*\* p<0.01, \*\*\* p<0.001).

	FN	TP
<i>Laboratory and Vitals (Mean Values)</i>		
SpO <sub>2</sub> /FiO <sub>2</sub> ***	472.9	202.7
Albumin *	3.7	3.3
GFR **	99.6	50.5
<i>Oxygen Device (N patients)</i>		
Room Air	7	3
Nasal Cannula	0	45
Mask	0	27

## References

1. Meiring, C. *et al.* Optimal intensive care outcome prediction over time using machine learning. *PLOS ONE* **13**, e0206862, DOI: [10.1371/journal.pone.0206862](https://doi.org/10.1371/journal.pone.0206862) (2018).
2. Kwon, J., Lee, Y., Lee, Y., Lee, S. & Park, J. An Algorithm Based on Deep Learning for Predicting In-Hospital Cardiac Arrest. *J. Am. Hear. Assoc.* **7**, DOI: [10.1161/JAHA.118.008678](https://doi.org/10.1161/JAHA.118.008678) (2018).
3. Komorowski, M., Celi, L. A., Badawi, O., Gordon, A. C. & Faisal, A. A. The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. *Nat. Medicine* **24**, 1716–1720, DOI: [10.1038/s41591-018-0213-5](https://doi.org/10.1038/s41591-018-0213-5) (2018).
4. Prosperi, M. C. F. *et al.* Investigation of expert rule bases, logistic regression, and non-linear machine learning techniques for predicting response to antiretroviral treatment. *Antivir. Ther.* **14**, 433–442 (2009).
5. Burdick, H. *et al.* Prediction of respiratory decompensation in covid-19 patients using machine learning: The ready trial. *Comput. Biol. Medicine* **124**, 103949, DOI: <https://doi.org/10.1016/j.combiomed.2020.103949> (2020).
6. Yadaw, A. S. *et al.* Clinical features of COVID-19 mortality: development and validation of a clinical prediction model. *The Lancet Digit. Heal.* **2**, e516–e525, DOI: [10.1016/S2589-7500\(20\)30217-X](https://doi.org/10.1016/S2589-7500(20)30217-X) (2020).
7. Li, W. T. *et al.* Using machine learning of clinical data to diagnose covid-19: a systematic review and meta-analysis. *BMC Med. Informatics Decis. Mak.* **20**, DOI: [10.1186/s12911-020-01266-z](https://doi.org/10.1186/s12911-020-01266-z) (2020).

**Table 4. Key feature differences between TP and FN groups predicted during epoch 2**

SpO<sub>2</sub>/FiO<sub>2</sub> measured in epochs 1 and 2, the type of oxygen delivery patients received in epochs 1 and 2, as well as the albumin and estimated glomerular filtration rate (GFR) test results from epoch 2 were found to be the most important for distinguishing patients for whom a false negative and true positive prediction was made at this stage of analysis. Feature importance was calculated via KL divergence and statistical testing for differences of means for laboratory and vital signs were conducted via ANOVA (\* p<.05, \*\* p<0.01, \*\*\* p<0.001).

	FN	TP
<i>Laboratory and Vitals (Mean Values)</i>		
SpO <sub>2</sub> /FiO <sub>2</sub> (From Epoch 1) ***	454.3	252.2
SpO <sub>2</sub> /FiO <sub>2</sub> (In Current Epoch) ***	477.9	231.5
Albumin (From Epoch 1) *	3.8	3.6
GFR (From Epoch 1) * *	84.4	46.8
<i>Oxygen Device in Epoch 1 (N patients)</i>		
Room Air	6	0
Nasal Cannula	1	16
<i>Oxygen Device in Epoch 2 (N patients)</i>		
Room Air	7	0
Nasal Cannula	0	13
Mask	0	3

8. Wongvibulsin, S. *et al.* Development of severe covid-19 adaptive risk predictor (scarp), a calculator to predict severe disease or death in hospitalized patients with covid-19. *Annals Intern. Medicine* DOI: [10.7326/M20-6754](https://doi.org/10.7326/M20-6754) (2021). <https://doi.org/10.7326/M20-6754>.
9. Obermeyer, Z. & Lee, T. H. Lost in thought — the limits of the human mind and the future of medicine. *New Engl. J. Medicine* **377**, 1209–1211, DOI: [10.1056/NEJMp1705348](https://doi.org/10.1056/NEJMp1705348) (2017). PMID: 28953443, <https://doi.org/10.1056/NEJMp1705348>.
10. Rajkomar, A., Dean, J. & Kohane, I. Machine learning in medicine. *New Engl. J. Medicine* **380**, 1347–1358, DOI: [10.1056/NEJMra1814259](https://doi.org/10.1056/NEJMra1814259) (2019). PMID: 30943338, <https://www.nejm.org/doi/pdf/10.1056/NEJMra1814259>.
11. Collins, F. S. & Varmus, H. A new initiative on precision medicine. *New Engl. J. Medicine* **372**, 793–795, DOI: [10.1056/NEJMp1500523](https://doi.org/10.1056/NEJMp1500523) (2015). PMID: 25635347, <https://doi.org/10.1056/NEJMp1500523>.
12. Johnson, A. E. W. *et al.* Machine learning and decision support in critical care. *Proc. IEEE. Inst. Electr. Electron. Eng.* **104**, 444–466, DOI: [10.1109/JPROC.2015.2501978](https://doi.org/10.1109/JPROC.2015.2501978) (2016).
13. Christie, S., Conroy, A., Callcut, R., Hubbard, A. & Cohen, M. Dynamic multi-outcome prediction after injury: Applying adaptive machine learning for precision medicine in trauma. *PLoS ONE* **14** (2019).
14. Zimmerman, J. E., Kramer, A. A., McNair, D. S. & Malila, F. M. Acute physiology and chronic health evaluation (apache) IV: Hospital: mortality assessment for today's critically ill patients. *Critical Care Medicine* **34**, 1297–1310, DOI: [10.1097/01.CCM.0000215112.84523.F0](https://doi.org/10.1097/01.CCM.0000215112.84523.F0) (2006).
15. Ghassemi, M. M., Mark, R. G. & Nemati, S. A visualization of evolving clinical sentiment using vector representations of clinical notes. In *2015 Computing in Cardiology Conference (CinC)*, 629–632, DOI: [10.1109/CIC.2015.7410989](https://doi.org/10.1109/CIC.2015.7410989) (IEEE, Nice, France, 2015).
16. Sharafoddini, A., Dubin, J. A. & Lee, J. Identifying subpopulations of septic patients: A temporal data-driven approach. *Comput. Biol. Medicine* **130**, 104182, DOI: [10.1016/j.combiomed.2020.104182](https://doi.org/10.1016/j.combiomed.2020.104182) (2021).
17. Ignatius, E. H. *et al.* Tocilizumab for the treatment of covid-19 among hospitalized patients: A matched retrospective cohort analysis. *Open Forum Infect. Dis.* **8**, 1–9, DOI: [10.1093/ofid/ofaa598](https://doi.org/10.1093/ofid/ofaa598) (2021). <https://doi.org/10.1093/ofid/ofaa598>.
18. Avery, R. K. *et al.* Inpatient covid-19 outcomes in solid organ transplant recipients compared to non-solid organ transplant patients: A retrospective cohort. *Am. J. Transplantation* **8**, 1–11, DOI: [10.1111/ajt.16431](https://doi.org/10.1111/ajt.16431) (2020). <https://doi.org/10.1111/ajt.16431>.
19. Boulanger, M. *et al.* Peripheral plasma cells associated with mortality benefit in severe covid-19: A marker of disease resolution. *The Am. J. Medicine* 1–5, DOI: [10.1111/ajt.16431](https://doi.org/10.1111/ajt.16431) (2021). <https://doi.org/10.1016/j.amjmed.2021.01.040>.

20. Garibaldi, B. T. *et al.* Comparison of time to clinical improvement with vs without remdesivir treatment in hospitalized patients with covid-19. *JAMA Netw. Open* DOI: [10.1001/jamanetworkopen.2021.3071](https://doi.org/10.1001/jamanetworkopen.2021.3071) (2021).
21. Pmap: The johns hopkins precision medicine analytics platform.
22. Karaba, S. M. *et al.* Prevalence of co-infection at the time of hospital admission in covid-19 patients, a multicenter study. *Open Forum Infect. Dis.* DOI: [10.1093/ofid/ofaa578](https://doi.org/10.1093/ofid/ofaa578) (2020).
23. Garibaldi, B. T. *et al.* Patient trajectories among persons hospitalized for covid-19. *Annals Intern. Medicine* DOI: [10.7326/m20-3905](https://doi.org/10.7326/m20-3905) (2020).
24. Andersen, K. M. *et al.* Association between chronic use of immunosuppressive drugs and clinical outcomes from coronavirus disease 2019 (covid-19) hospitalization: A retrospective cohort study in a large us health system. *Clin. Infect. Dis.* DOI: [10.1093/cid/ciaa1488](https://doi.org/10.1093/cid/ciaa1488) (2021).
25. Iwendi, C. *et al.* Covid-19 patient health prediction using boosted random forest algorithm. *Front. Public Heal.* **8**, 357, DOI: [10.3389/fpubh.2020.00357](https://doi.org/10.3389/fpubh.2020.00357) (2020).
26. McInnes, L., Healy, J. & Melville, J. Umap: Uniform manifold approximation and projection for dimension reduction (2020). [1802.03426](https://arxiv.org/abs/1802.03426).
27. Joe H. Ward, J. & Hook, M. E. Application of an hierarchical grouping procedure to a problem of grouping profiles. *Educ. Psychol. Meas.* **23**, 69–81, DOI: [10.1177/001316446302300107](https://doi.org/10.1177/001316446302300107) (1963). <https://doi.org/10.1177/001316446302300107>.
28. Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
29. Caliński, T. & Harabasz, J. A dendrite method for cluster analysis. *Commun. Stat.* **3**, 1–27, DOI: [10.1080/03610927408827101](https://doi.org/10.1080/03610927408827101) (1974). <https://www.tandfonline.com/doi/pdf/10.1080/03610927408827101>.
30. Kullback, S. & Leibler, R. A. On Information and Sufficiency. *The Annals Math. Stat.* **22**, 79 – 86, DOI: [10.1214/aoms/1177729694](https://doi.org/10.1214/aoms/1177729694) (1951).
31. Cui, W., Robins, D. & Finkelstein, J. Unsupervised machine learning for the discovery of latent clusters in covid-19 patients using electronic health records. *Stud. health technology informatics* **272**, 1—4, DOI: [10.3233/shti200478](https://doi.org/10.3233/shti200478) (2020).
32. Matelsky, J. *et al.* Container-based clinical solutions for portable and reproducible image analysis. *J. Digit. Imaging* **31**, 315–320, DOI: [10.1007/s10278-018-0089-4](https://doi.org/10.1007/s10278-018-0089-4) (2018).