

# Clustering US Counties to Find Patterns Related to the COVID-19 Pandemic

Cora Brown  
Sarah Milstein  
Tianyi Sun  
Cooper Zhao

University of Minnesota, School of Mathematics  
in collaboration with Ecolab  
November 16, 2020

# Table of Contents

1	Introduction	2
2	Data	2
3	Evaluation Metrics and Interpretation Techniques	4
3.1	Evaluation Metrics . . . . .	4
3.2	Interpretation Techniques . . . . .	5
4	Clustering Techniques	6
4.1	$K$ -Means Clustering . . . . .	6
4.2	Mini Batch $K$ -Means Clustering . . . . .	6
4.3	Fuzzy $c$ -Means Clustering . . . . .	7
4.4	Gaussian Mixture Models . . . . .	7
4.5	Hierarchical Clustering . . . . .	7
4.6	OPTICS Clustering . . . . .	8
4.6.1	DBSCAN . . . . .	8
4.6.2	OPTICS . . . . .	8
5	Results	9
5.1	$K$ -means results . . . . .	9
5.2	Fuzzy- $c$ means . . . . .	12
5.3	Gaussian Mixture Model . . . . .	14
5.4	Mini Batch $K$ -Means . . . . .	16
5.5	Hierarchical Clustering . . . . .	17
5.5.1	Distinguishing Cluster 0 from All Other Clusters . . . . .	18
5.5.2	Distinguishing the Smaller Clusters . . . . .	20
5.5.3	Summary of Hierarchical Clustering Results . . . . .	22
5.6	OPTICS Clustering . . . . .	22
5.6.1	Dimension Reduction and OPTICS Clustering . . . . .	22
5.6.2	Summary of OPTICS Results . . . . .	23
6	Conclusion	26
	References	27

# 1 Introduction

When COVID-19 first started spreading and quarantine was implemented, the Society for Industrial and Applied Mathematics (SIAM) Student Chapter at the University of Minnesota-Twin Cities began a collaboration with Ecolab to use our skills as data scientists and mathematicians to extract useful insights from relevant data relating to the pandemic. This collaboration consisted of multiple groups working on different projects. In this write-up we focus on using clustering techniques to help us find groups of similar counties in the US and use that to help us understand the pandemic. Our team for this project consisted of University of Minnesota students Cora Brown, Sarah Milstein, Tianyi Sun, and Cooper Zhao, with help from Ecolab Data Scientist Jimmy Broomfield and University of Minnesota student Skye Ke.

In the sections below we describe all of the work done for this project. In Section 2, we list the data we gathered, as well as the feature engineering we performed. In Section 3, we describe the metrics we used for evaluating our models. In Section 4, we explain the methods we used for interpreting the results of our various clustering approaches. In Section 5, we describe the different clustering methods we implemented. In Section 6, we present the results of our clustering techniques and provide relevant interpretation. Finally, in Section 7, we provide some concluding remarks comparing the different clustering methods.

## 2 Data

For this project, we used data from a variety of sources, including the collection of data for the UNCOVER COVID-19 Kaggle Challenge available at <https://www.kaggle.com/roche-data-science-coalition/uncover> and case and death numbers from the Centers for Disease Control.

Out of the available data for the UNCOVER Challenge, we chose and engineered features which could be relevant in COVID-19 clustering. Here we group our final set of features into several groups in order to more easily describe their relevance to the problem at hand. First, we have features which represent the makeup and demographics of each county, including

1. Area - square miles of area in each county,
2. Population - population of each county,
3. Ranking: socioeconomic status - To determine this ranking, the CDC first computed the percentile ranking (0 to 1 with 1 meaning greater vulnerability) for the following features, based on Census data, that they determined to be in the category of socioeconomic status: (1) percentage of people below the poverty line, (2) percentage of people, age 16 and older, who are unemployed, (3) per capita income, and (4) percentage of people, age 25 and older, without a high school diploma. They then summed together the percentile ranking across all 4 features for each county, ranked these resulting sums, and computed each county's percentile ranking with respect to this sum as the final column value.
4. Ranking: household decomposition and disability - Computed following the same procedure as 'ranking: socioeconomic status' except the features used were: (1) percentage of people aged 65 and older, (2) percentage of people aged 17 and younger, (3) percentage of civilians with a disability, (4) percentage of single-parent households with children under 18.

5. Ranking: minority status and language - Computed following the same procedure as ‘ranking: socioeconomic status’ except the features used were: (1) percentage of people who are not white, non-Hispanic, and (2) percentage of people, age 5 and older, who self-reported speaking English ‘less than well’.
6. Ranking: housing and transportation - Computed following the same procedure as ‘ranking: socioeconomic status’ except the features used were: (1) percentage of housing in structures with more than 10 units, (2) percentage of mobile homes, (3) percentage of occupied housing with more people than rooms, (4) percentage of households with no vehicle available, and (5) percentage of people in institutionalized group quarters.
7. Index of relative rurality - measure of the level of rurality of a county based on population size, density, remoteness, and built-up area, where 0 indicates that a county has a low level of rurality, i.e. urban, and 1 indicates that a county is the most rural.
8. ICU beds - number of ICU beds available in each county,
9. Nursing home population - the nursing home population in each county.

Next, we have COVID-19 specific data, including

1. Number of testing locations - testing locations available in each county,
2. Google mobility score - score representing the amount movement (i.e. daily travel) was decreased during state or county shutdowns,
3. State closure status - the status of state shutdowns,
4. School closure status - the status of school closures.

This second set of features represents both the state and county’s policies with respect to the COVID-19 pandemic through testing locations and state and school closure statuses, and it represents the public’s response to the pandemic in each county through the Google mobility score. This mobility score is an aggregate measure of the amount that movement between locations such as homes and grocery stores changed due to the pandemic. A county in which the population severely limited their movements at the beginning of the pandemic will have a high mobility score, while a county which remained active will have a low score.

The final group of features, listed below, represents quantities that were engineered from the time-series of cumulative confirmed cases and deaths in each county, from January 22 to August 8 (the day of data collection). Note that January 22 is around the date of the first confirmed COVID-19 case in the US. In order to make the clustering easier, we took the time-series data and described it with a series of “summary” features which attempt to capture important aspects of the patterns discernible in the time-series. These features include:

1. Case growth rate to first peak - daily rate of change of case numbers from January 22 to April 12 (around the first peak of COVID-19 cases in the US); this metric represents the growth rate of case numbers in each county early in the pandemic
2. Last month case growth rate - daily rate of change of case numbers from July 8 to August 8; representing the growth rate of case numbers in each county later in the pandemic
3. New cases April 12 - number of new confirmed cases reported in each county on April 12,

4. Deaths April 12 - number of new confirmed deaths reported in each county on April 12,
5. New cases July 23 - number of new confirmed cases reported in each county on July 23 (around the second peak of COVID-19 cases in the US),
6. Deaths July 23 - number of new confirmed cases reported in each county on July 23,
7. Cumulative cases August 8 - number of total cumulative confirmed cases in each county on August 8 (the date of data collection),
8. Cumulative deaths August 8 - number of total cumulative deaths in each county on August 8 (the date of data collection).

In addition to the new case/death numbers on April 12 (around the first USA-wide peak) and July 23 (around the second USA-wide peak), we also included growth rates of the new case values early in the pandemic and in the month leading up to August 8, and the total cumulative cases/deaths in each county up to August 8. These features aim to capture how the number of confirmed COVID-19 cases and deaths were changing throughout the summer, and also what the raw values looked like during the US case peaks.

### 3 Evaluation Metrics and Interpretation Techniques

Below we describe various methods that we used both for evaluating the clustering methods that we implemented and for interpreting the results of these methods. The content of this section comes from [5] and [6].

#### 3.1 Evaluation Metrics

In order to determine which clustering techniques perform well on our data set, we outline several evaluation metrics. While these evaluation metrics cannot offer a guarantee of accuracy since there is no definitive ground truth, they allow us to determine whether or not the clusters we see might represent meaningfully distinct groups of counties. In addition, these evaluation methods can help tune hyperparameters when necessary (e.g. determining the number of clusters for  $K$ -Means clustering). Below we list some of the metrics we used.

- The *silhouette score* gives the average over all points in our dataset of how similar each datapoint is to its own clusters compared to the next closest cluster. Here, the values range from  $-1$ , indicating clusters that do not separate the datapoints well, to  $1$ , indicating clusters that clearly distinguish the datapoints.
- The *Calinski-Harabasz Index* is defined as the ratio of the sum of intra-cluster dispersion and inter-cluster dispersion for all clusters, where dispersion is a measure of distance. The index is high when clusters are dense and well-separated.
- The *Davies-Bouldin Index* indicates the average “similarity” between clusters by comparing the distance between clusters with the size of the clusters. Here,  $0$  is the smallest possible score, and values closer to  $0$  represent better clusterings.
- The *elbow method*, also called *distortion*, is a heuristic used to help determine the number of clusters to use for a particular data set. This method consists of plotting the explained

variation as a function of the number of clusters, and picking the “elbow” of this resulting curve as the number of clusters to use.

- The *Bayesian information criterion (BIC)* is a criterion for model selection among a finite set of models, specifically for Gaussian mixtures parameter selection. The BIC attempts to choose the model with the highest likelihood while penalizing models with more parameters. The model with the lowest BIC is preferred.
- The *Akaike information criterion (AIC)* estimates the information lost by a given model. The less information a model loses, the higher the quality of that model. Similar to the BIC, the AIC tries to find the model with the highest likelihood while penalizing for adding more parameters to the model. However, the penalty for adding more parameters is lower for the AIC than for the BIC.

### 3.2 Interpretation Techniques

Below, we list a few methods we used to interpret the results of applying a particular clustering method. While these methods will not necessarily tell us about how the clusters were created, they can help us understand what makes counties in one cluster similar and what makes two clusters different.

- We can find which features have values that differ the most between cluster. That is, if the average values of a feature differ significantly between clusters, that variable can help us distinguish those clusters.
- We can extract a feature importance score from building a random forest, and then select the features with the highest scores. For this approach, we set the cluster labels as the target value and build a random forest. Then, based on the resulting random forest, we can extract an importance score for each feature. We can interpret the features with the highest scores and features that can help distinguish our clusters.
- We can describe our clusters or groups of clusters using a decision tree. Here, we build a decision tree using the cluster labels as the output variable. Then, we can see which clusters were grouped near by based on the decision tree, and use the feature value splits in the tree to interpret the clusters or groups of clusters.
- To see if our clusters represent natural breaks in particular features, we can use the Jenks Optimization Method to cluster using each feature individually. Jenks Optimization Method does single variable clustering. This method uses an iterative approach to group data by gradually minimizing variance within classes and maximizing variance between classes, which is similar to  $K$ -Means clustering on multiple features. Then, we can use the  $v$  measure score to see how similar the clustering from the Jenks Optimization method is to our original clustering, generated by using all features. A higher  $v$  measure score indicates a higher agreement of two clustering approaches, which would suggest that the clusters we found originally represent natural breaks in each feature used for the Jenks Optimization method.

## 4 Clustering Techniques

In this section, we describe the various clustering methods that we implemented in our work. Since one motivation for this project was to learn about different clustering techniques, we tried a wide range of approaches. These include  $K$ -means and other prototype based methods (Sections 4.1, 4.2, 4.3, 4.4), hierarchical clustering (Section 4.5), and density-based methods (Section 4.6). The content of these sections comes from [5] and [6]. We also use [2] for the  $K$ -means and other prototype-based approaches, and we use [1], [3], and [4] for the section on OPTICS.

### 4.1 $K$ -Means Clustering

The  $K$ -means algorithm clusters data by separating samples into  $K$  groups of equal variance. The algorithm divides a set of samples into disjoint clusters, each described by the mean of the samples in the cluster, called the “centroid.” Note that the centroid may not be an actual sample in the cluster. The  $K$ -means algorithm has the following four steps:

1. Randomly initialize  $K$  samples from the dataset as centroids.
2. Assign the rest of the samples in the dataset to their nearest centroid.
3. Update the centroids by taking the mean value of all the samples in each cluster.
4. Repeat steps 2 and 3 until the clusters stop changing.

For this algorithm,  $K$  must be specified before starting this algorithm. In the results section below, we describe the method we used to select an appropriate value for  $K$ .

We note that  $K$ -means performs poorly on elongated clusters because within-cluster sum-of-squares is small only when clusters are grouped densely. In addition, within-cluster sum-of-squares can easily be large in high-dimensional spaces since datapoints are often more spread out. In order to address this latter issue, we can perform Principle Component Analysis (PCA) in order to reduce the dimensions of the sample space prior to using  $K$ -means clustering.

### 4.2 Mini Batch $K$ -Means Clustering

Mini Batch  $K$ -Means is a variant of the  $K$ -means algorithm with the same first step but where we update the centroids using only some of the data at a time. This algorithm runs as follows:

1. Randomly initialize  $K$  samples from the dataset as centroids.
2. Draw  $b$  samples randomly from the dataset to form a mini-batch, and assign each sample to the nearest centroid.
3. Update the centroid of each cluster by taking the average of the newly added samples from that mini-batch as well as all the previous samples assigned to that centroid.
4. Repeat steps 2 and 3 until the clusters stop changing.

The results of Mini Batch  $K$ -Means are often similar to those of  $K$ -means, but the mini batch version of the algorithm converges faster.

### 4.3 Fuzzy $c$ -Means Clustering

Fuzzy clustering is a form of clustering in which each data point can belong to more than one cluster. Fuzzy  $c$ -means is one of the most common such methods. It is similar to the  $K$ -means algorithm, but it also involves a coefficient value for each pair of a datapoint and a cluster, representing the degree to which that datapoint falls into that cluster. Fuzzy  $c$ -means uses the following steps:

1. Choose a number of clusters  $c$ .
2. Randomly initialize coefficients for each sample for being in each cluster.
3. Compute the centroid for each cluster by averaging over all datapoints, where we weight each datapoint by its coefficient.
4. Compute the new coefficients for each datapoint-cluster pair.
5. Repeat steps 3 and 4 until the algorithm converges.

### 4.4 Gaussian Mixture Models

A Gaussian mixture model is a probabilistic model that assumes our datapoints are generated from a mixture of a finite number, say  $k$ , of Gaussian distributions with unknown parameters. Each Gaussian distribution in the mixture is composed of the following parameters:

- A mean  $\mu_k$  defining the center of the distribution,
- A covariance matrix  $\Sigma_k$  defining the variance of the distribution, and
- A mixing probability  $\pi_k$  defining how big the Gaussian function will be, or the relative sizes of the different Gaussians.

We can then find these parameters by iteratively updating estimations for them. In Gaussian mixture model selection process, we also need to specify the covariance type and the number of Gaussians to use. We can use BIC and AIC to help us select both of these parameters.

### 4.5 Hierarchical Clustering

One common method for finding structure in unlabeled data is hierarchical clustering. Unlike many other clustering approaches, this technique does not find one specific number of clusters within a dataset, but instead finds all numbers of clusters at once. Then, from the resulting hierarchy, we can select the number of clusters that we want to use.

There are two main types of hierarchical clustering: agglomerative and divisive. Agglomerative clustering starts with all observations in their own cluster, and sequentially joins the two closest clusters until every observation is in one cluster together. Divisive clustering, on the other hand, starts with a single cluster of all the observations and sequentially splits off one observation. Here, we focus on the agglomerative approach.

When applying agglomerative clustering, we need to specify (1) a distance metric between observations and (2) a definition of distance between clusters, rather than just individuals. In our



work on this project, we used `linkage` from `scipy.cluster.hierarchy`, which has a number of available distance metrics, and cluster distances, also called linkages. For distance metric, we tested almost all that are available in `pdist` from `scipy.spatial.distance`, including Euclidean, Minkowski, cityblock (or Manhattan), cosine, and squared Euclidean distances. For distance between clusters, we tested all available options, including Ward's, complete, average, and single linkages. We also specified a final desired number of clusters, so that as we are iteratively merging clusters, we stop once we have this number remaining.

## 4.6 OPTICS Clustering

In this section, we will describe the OPTICS clustering algorithm. First, though, we will introduce the DBSCAN clustering method.

### 4.6.1 DBSCAN

Density-based spatial clustering of applications with noise (DBSCAN) is a density-based non-parametric algorithm which locates regions of high density in the data that are separated from one another by regions of low-density. There are several approaches to defining the density of the data points. In this section, we will mainly use the center-based approach for defining the density.

In the center-based approach, we have `MinPts` and `Eps`, both of which are user-specified parameters representing the number of points and the radius. Then, we classify all the data points in the given space into one of the following three categories:

- **Core Points:** We label a datapoint as a core point if there are at least `MinPts` datapoints (including the point itself) within a distance of radius `Eps` of this point.
- **Border Points:** we label a datapoint as a border point if it is in the neighborhood of a core point. That is, it is within `Eps` of a core points, but there are fewer than `MinPts` datapoints within a distance of `Eps` from this border point.
- **Noise:** We label a datapoint as noise if it is neither a core point nor a border point.

After labeling each point, we then form our clusters. We group two core points into one cluster if the distance between them is less than `Eps`. We group a border point into a cluster if the border point is within a distance of `Eps` from a core point in that cluster. Finally, we group all the noise points together into one cluster. The resulting clusters are the output of the DBSCAN algorithm.

### 4.6.2 OPTICS

We notice that DBSCAN uses a fixed distance threshold (`Eps`) which can not detect clusters of varying densities, which can sometimes lead to poor clustering. Ordering Points To Identify the Clustering Structure (OPTICS) has a similar algorithm to DBSCAN. However, OPTICS addresses this problem of a fixed distance by using a steepness threshold for `Eps`.

In the OPTICS algorithm as in DBSCAN, we specify a value `MinPts`, which gives us a lower bound on the number points in the neighborhood of some point in order for us to consider that

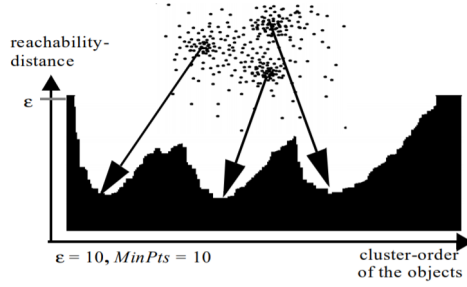


Figure 1: OPTICS Reachability Graph Example [1]

point, and a value  $\text{Eps}$ , which is the maximum radius of a neighborhood around a point that we will consider. Then in the OPTICS algorithm, for each datapoint  $p$ , we determine its core distance, which is the minimum radius size so that there are  $\text{MinPts}$  points within that distance. Note that if this radius is larger than  $\text{Eps}$ , we set the core distance of  $p$  to be undefined. Then for any point  $q$  within the  $\text{Eps}$ -neighborhood of  $p$ , we set the reachability distance from  $p$  to  $q$  to be the maximum of the core distance of  $p$  and of the distance from  $p$  to  $q$ . If  $q$  is further than  $\text{Eps}$  distance from  $p$ , we let the reachability distance from  $p$  to  $q$  be undefined.

Next, we order our datapoints so that points near each other with respect to reachability distance are near each other in the ordering. This produces a reachability graph as in Figure 1, where reachability score is plotted against the datapoint ordering. The spikes in this graph of high reachability scores correspond to less densely connected pointed, and thus indicate potential boundaries between clusters. Then, we can set different reachability thresholds and set the resulting clusters to be all points between spikes of at scores higher than the threshold.

To apply the OPTICS Clustering to our data, we used the package from `sklearn.cluster.OPTICS` which requires two parameters: `min_samples` ( $\text{Mps}$ ) and `metric`. To select these parameters, we implemented a grid search over a range of 2-30 for `min_samples` and over almost all the metrics in `scikit-learn` and `scipy.spatial.distance`, and selected the best pair of parameters using the silhouette score and Calinski-Harabasz index.

## 5 Results

In the following sections, we describe our results from implementing the clustering algorithms described in Section 4 to the data we gathered described in Section 2.

### 5.1 $K$ -means results

Since the  $K$ -means clustering method does not work well with high-dimensional data, as mentioned in Section 4.1, we first used principle component analysis (PCA) to visualize the data and reduce dimensionality before clustering. This process showed that by using 8 principle components, we were able to capture approximately 95% of the variance in our data. This preprocessing step reduced our feature space by more than half.

Next, we chose  $K$ , the number of clusters, using the elbow score, silhouette score, and Calinski-Harabasz index. We used the package `yellowbrick` to calculate and visualize these scores since it has an indicator that automatically shows the optimal number of components based on each evaluation methods that balances a high score with using more clusters. The optimal number of

components according to the elbow score is 8. The Calinski-Harabasz index has a similar shape to the elbow score plot, which decreases quickly between 4 and 7 clusters, and then decreases more slowly afterwards. This suggests that perhaps around 7 or 8 clusters is reasonable, as well. However, the silhouette score drops off sharply between 2 and 4 clusters and then stays around the same score through at least 18 clusters. We can see the relevant plots of these scores with the corresponding best number of features indicated in Figure 2.

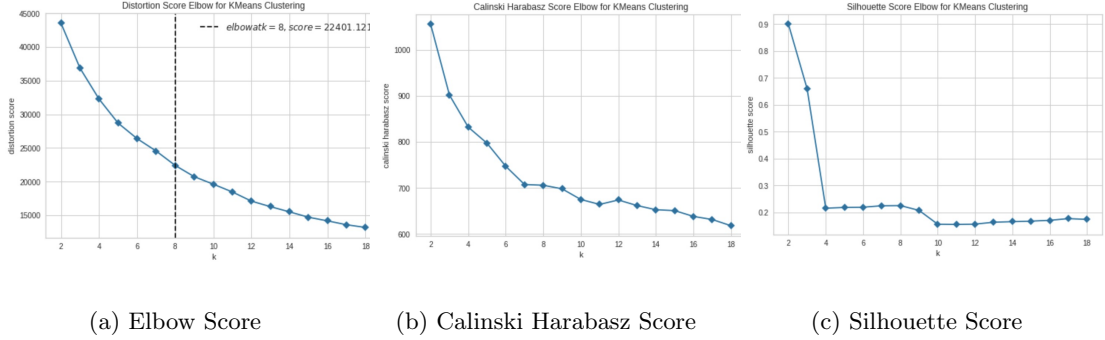


Figure 2: K-Means Evaluation Scores

Because both elbow score and Calinski-Harabasz index both suggested that around 8 clusters was a good choice, we choose this the number of clusters to use. After performing  $K$ -means clustering with 8 clusters, we can visualize the resulting clusters in the map in Figure 3, where the color of a county corresponds to the cluster that it is in. We note that many counties in the southern half of the United States, as well as those in the west appear to be clustered together, and many counties in the Midwest as well as the Northeast appear to be clustered together.

To interpret the clusters, first, we selected the important features using random forest feature selection as described in Section 3.2. The top seven important features are listed here:

- county ranking with respect to socioeconomic status,
- county ranking with respect to household decomposition and disability,
- county ranking with respect to minority status and language,
- county population,
- cumulative cases as of August 8,
- number of ICU beds.

Figure 4 shows the average variation of each features among different clusters.

We also used Jenks Optimization Method to cluster each features individually, and then we used the  $v$ -measure score between each feature clusters and the overall clusters to test the homogeneity between these sets of clusters. A high score indicates a high homogeneity. The features with the highest scores for  $K$ -Means clustering are: cumulative cases as of August 8, county ranking with respect to socioeconomic status, and county population. This suggests that natural breaks in the values for these features match up with to the results of  $K$ -means clustering, meaning that our clusters seem to correspond to different values of these features.

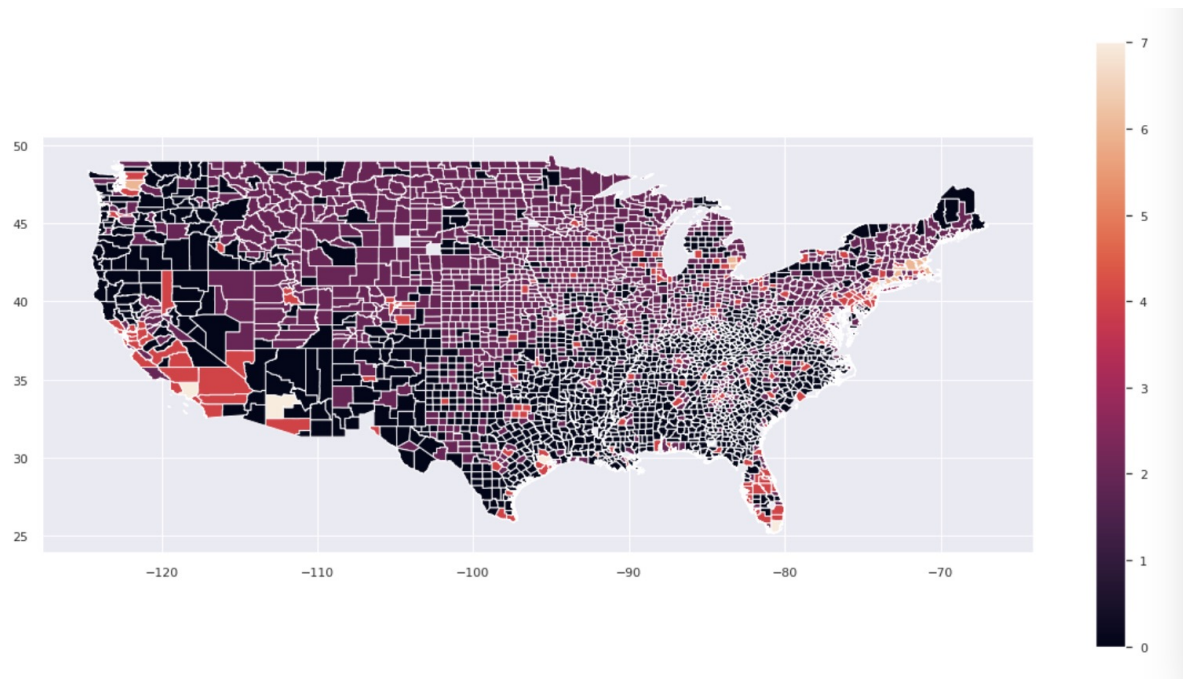


Figure 3: K-Means Map

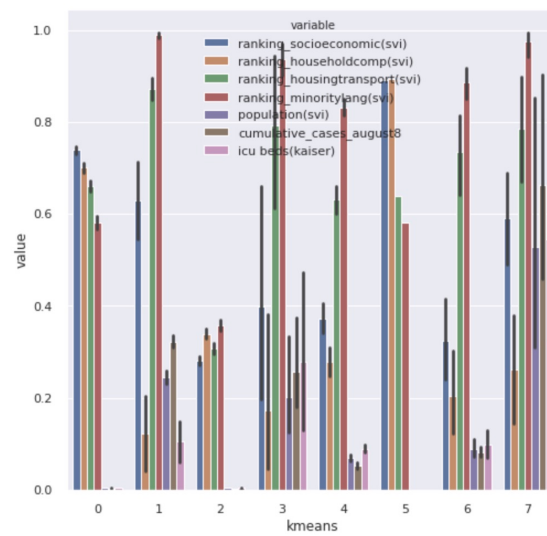
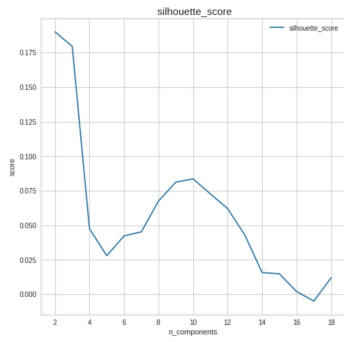
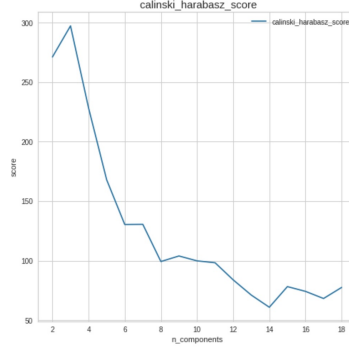


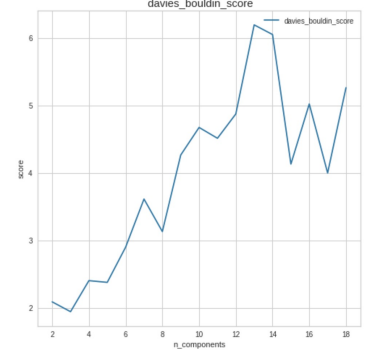
Figure 4: K-Means clusters Important Features



(a) Silhouette Score



(b) Calinski Harabasz Score



(c) Davies Bouldin Score

Figure 5: Fuzzy- $c$  Means Evaluation Scores

## 5.2 Fuzzy- $c$ means

To select the optimal number of clusters for Fuzzy- $c$  means clustering, we used the silhouette score, the Calinski-Harabasz score, and the Davies-Bouldin score similarly to in the section above. As we can see in Figure 5, each of these returned 3 as the optimal number of clusters.

After performing Fuzzy- $c$  means clustering with 3 clusters, we can visualize the resulting clusters in the map in Figure 6, where the color of a county corresponds to the cluster that it is in. Note that the clusters here appear to be very similar to the clusters we found using  $K$ -means clustering. This makes sense since these two methods are quite similar.

The top seven important features according to random forest feature selection are listed here:

- county ranking with respect to socioeconomic status,
- county ranking with respect to housing and transportation,
- county ranking with respect to household decomposition and disability,
- county ranking with respect to minority status and language,
- county population,
- index of relative rurality,
- cumulative cases as of August 8.

Figure 7 shows the average variation of each features among different clusters. The results show that the values for ranking with respect to socioeconomic status, ranking with respect to housing and transportation, and ranking with respect to household decomposition differ between the clusters. Therefore, we are able to use them distinguish the clusters we got from Fuzzy- $c$  Means clustering.

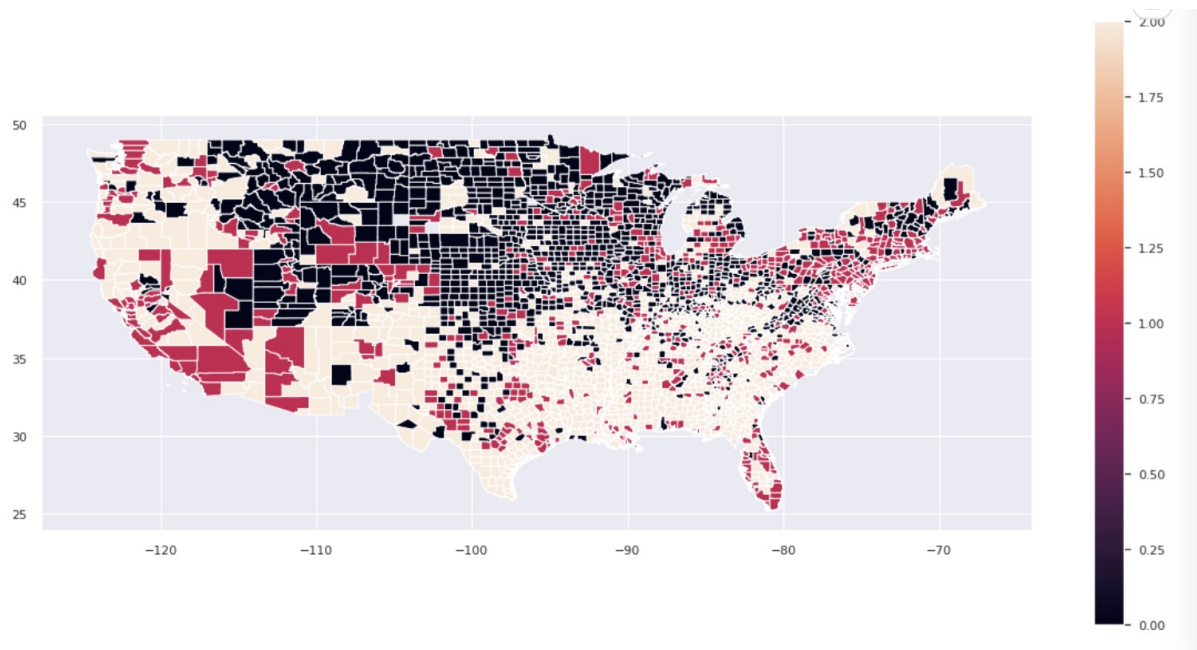


Figure 6: Fuzzy-c Means Map

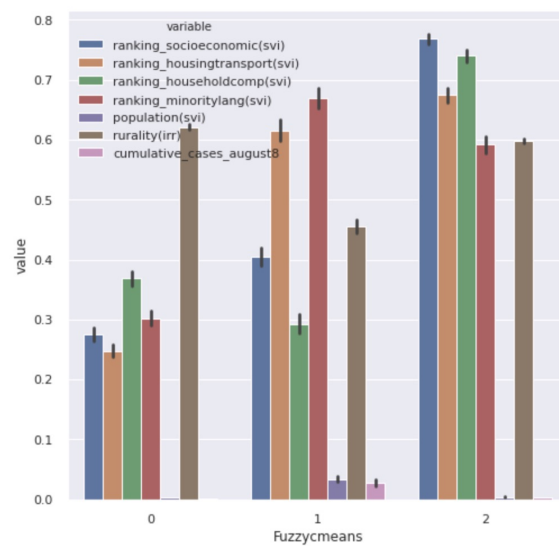


Figure 7: Fuzzy-c Means Clusters Important Features

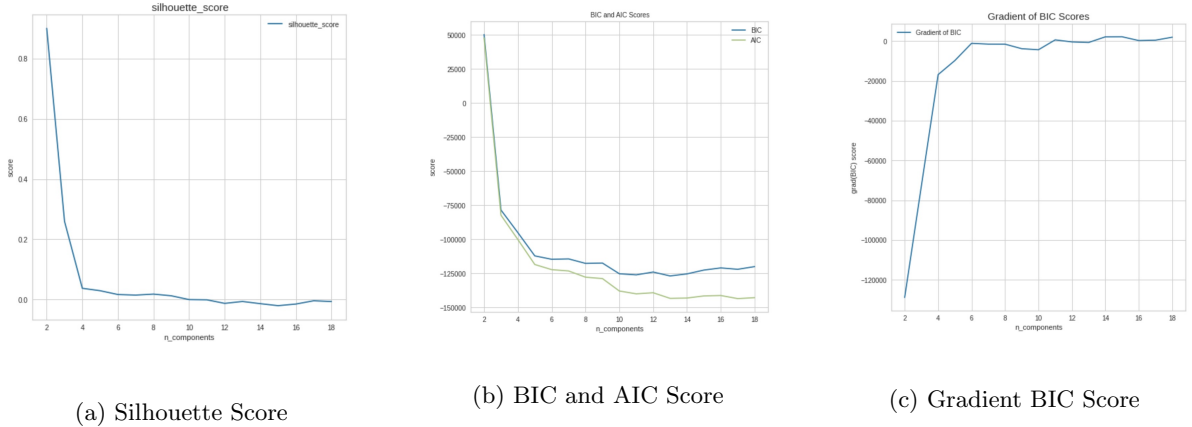


Figure 8: Gaussian Mixture Model Evaluation Scores

### 5.3 Gaussian Mixture Model

For the Gaussian Mixture model approach, the evaluation methods we used were silhouette score, BIC, and AIC. BIC and AIC both return an optimal number of clusters of either 3 or 5, but the silhouette score is higher for 3 clusters, meaning the best choice is 3 clusters, as we can see in Figure 8.

After building a Gaussian Mixture model with 3 clusters, we can visualize the resulting clusters in the map in Figure 9, where the color of a county corresponds to the cluster that it is in. Here we see that almost every county is in cluster 0 or cluster 1, and very few counties are in cluster 2.

The top seven important features according to random forest feature selection are listed here:

- new cases on July 23,
- case growth rate to first peak,
- county population,
- cumulative cases as of August 8,
- cumulative deaths as of August 8,
- number of ICU beds,
- index of relative rurality.

Figure 10 shows the average variation of each feature across the different clusters. This shows that the average value for index of relative rurality differs greatly between the clusters we got from the Gaussian mixture model. Therefore, we can use this feature to distinguish our resulting clusters, where on average cluster 0 contains the more rural counties, cluster 1 contains the less rural counties, and cluster 2 contains counties somewhere in the middle with regards to rurality.

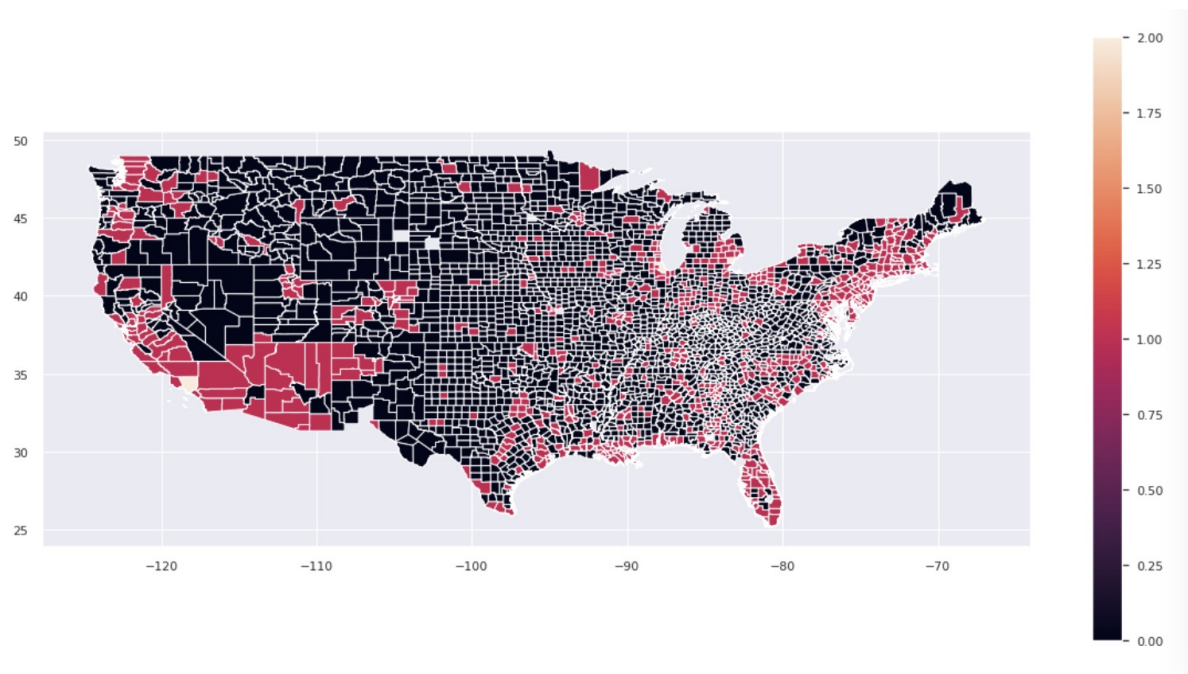


Figure 9: Gaussian Mixture Map

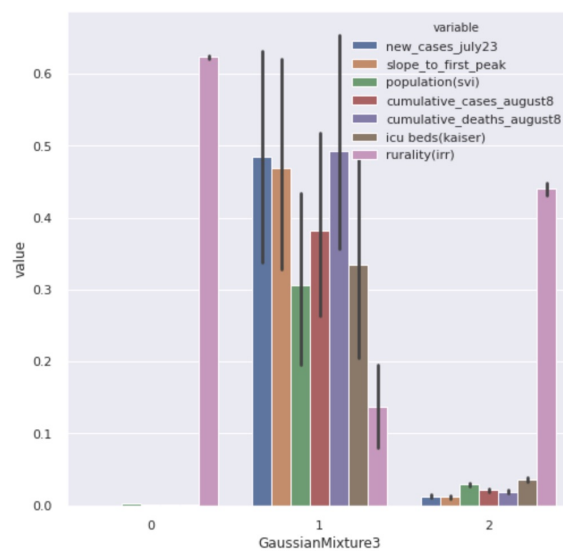
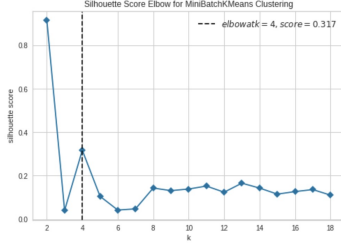
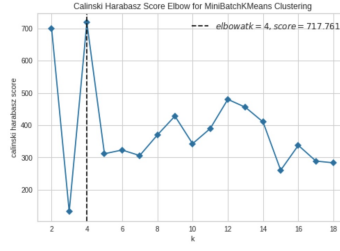


Figure 10: Gaussian Mixture Important Features

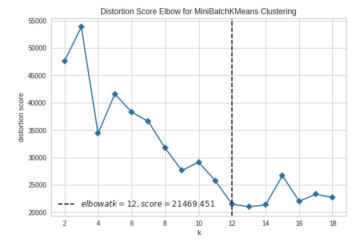




(a) Silhouette Score



(b) Calinski-Harabasz Score



(c) Gradient Elbow Score

Figure 11: Mini Batch  $K$ -Means Evaluation Scores

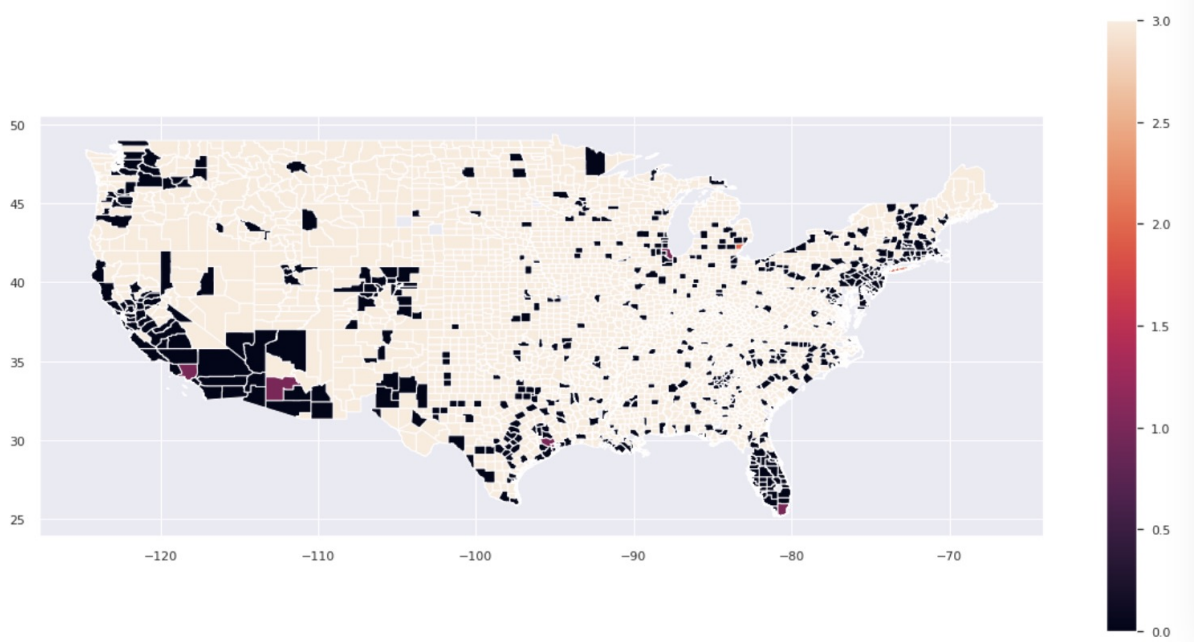


Figure 12: Mini Batch  $K$ -Means Map

#### 5.4 Mini Batch $K$ -Means

Now, we describe our results from applying Mini Batch  $K$ -means. In order to determine the number of clusters for this approach, we used the evaluation methods silhouette score, Calinski-Harabasz score, and elbow score. Both silhouette score and Calinski-Harabasz score indicate that 4 is a good choice for the number of clusters, as we can see in Figure 11. The elbow score indicates that 12 is a good choice for the number of clusters. However, since the corresponding elbow score with 4 clusters is a local minimum, we decided to use 4 clusters for Mini Batch  $K$ -means.

After building a Mini Batch  $K$ -means model with 4 clusters, we visualized the resulting clusters in the map in Figure 12, where the color of a county corresponds to the cluster that it is in.

The top seven important features according to random forest feature selection are listed here:

- county population,

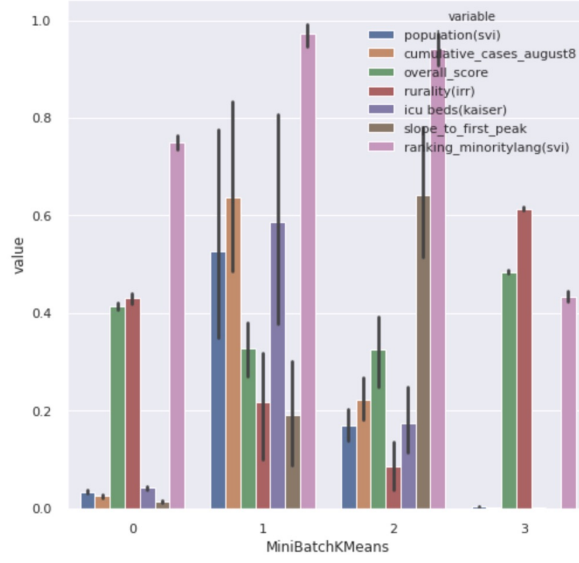


Figure 13: Mini Batch  $K$ -Means Important Features

- cumulative cases as of August 8,
- mobility score,
- index of relative rurality,
- number of ICU beds,
- case growth rate to first peak,
- county ranking with respect to minority status and language.

Figure 13 shows the average variation of each features among different clusters. This shows that we can distinguish our clusters by considering different values for the features: county population, cumulative cases as of August 8, index of relative rurality, and number of ICU beds.

## 5.5 Hierarchical Clustering

To find the best combination of linkage, distance metric, and number of clusters for hierarchical clustering, we trained a model with each different set of parameters on our dataset, computed its corresponding silhouette score, and returned a set of parameters with balance of a high score and a “reasonable” number of clusters. We note that the highest score always went to the clustering with as few clusters as possible. To counteract this, we chose the set of parameters that had the largest number of features that also had a silhouette score above 0.5.

Before running this grid search, we applied Principal Component Analysis (PCA) to our data to reduce the number of features in our dataset. PCA works by replacing the features in our dataset with new features that are linear combinations of the original features. It aims to capture the directions of greatest variance in the data, and reduce any redundancy across our feature space. Here, we reduced the original set of features down to 8 principal component features, which together capture 90% of the variance we see in our data. Applying our grid search method to

this feature-engineered dataset and following the process described in the previous paragraph, we chose the following set of parameters: average linkage (the distance between clusters is the average of all pairwise distances), Euclidean distance metric, and 25 clusters.

In Figure 14 we see a graph of the counties in our dataset, excluding those in HI and AK, colored by their cluster number. As can be seen in the graph, there is one cluster, which we refer to as Cluster 0, that contains most of the counties. In particular, it contains 2971 counties, which is almost 99% of all the counties in our dataset. Of the remaining 24 clusters, 13 contained only 1 county, and the remaining 11 clusters contained between 2 and 31 counties, for a total of 106 counties outside the one large cluster.

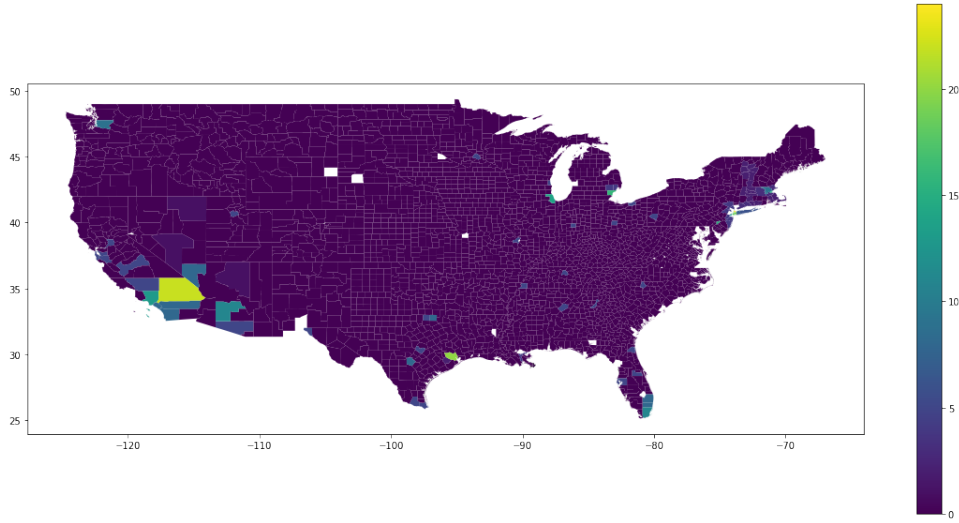


Figure 14: Map of US Counties, colored by hierarchical cluster.

In the following few sections, we present some insights from this hierarchical clustering approach. First, we describe the characteristics that distinguish counties in Cluster 0 from the remaining counties. Then, we discuss what distinguishes the clusters of counties outside Cluster 0 from each other. Finally, we present some comments on the overall results of hierarchical clustering applied to this dataset.

### 5.5.1 Distinguishing Cluster 0 from All Other Clusters

We first explore the differences between the counties in Cluster 0 and the counties outside that one cluster, by building a decision tree classifier to distinguish these two groups, shown in Figure 15. The first split in the resulting tree looks at whether a county has more or less than 12,626 cumulative COVID-19 cases as of August 8, which is more than 1.5 standard deviations above the average value of about 1600. Of the counties with more than 12,626 cumulative cases, we see that 65 of them are not in Cluster 0, which is about 61% of the remaining counties, and 10 of them are in Cluster 0, which is only about 0.3% of the counties in that cluster. Thus, we note that over half of the non-Cluster 0 counties have an incredibly high cumulative COVID-19 case count, and make up a majority of counties with such high cumulative case count.

Along the left branch of our decision tree, for which counties had fewer than 12,626 cumulative COVID-19 cases as of August 8, the next split is whether a county has more or less than 10 COVID-19 testing locations, which is about 3 standard deviations above the mean of 1. Of the counties with more than 10 testing locations, we have 22 counties outside Cluster 0, which

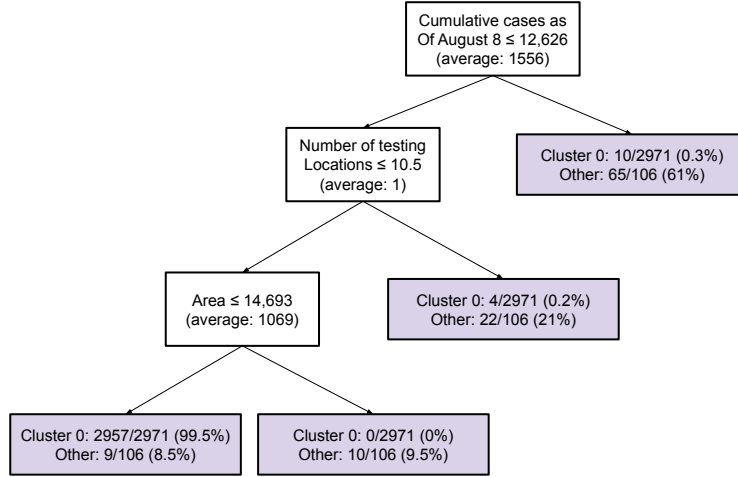


Figure 15: Simplified Decision Tree Distinguishing Cluster 0 from All Other Clusters.

makes up another 21% of such counties, and we have 4 counties in Cluster 0, which is only about 0.2% of Cluster 0 counties. Thus, we see that another subset of the non-Cluster 0 counties are those with an especially large number of testing locations.

The last split we point out from this tree is again along the left branch of the previous split, for which counties had fewer than 10 testing locations. This split looks at whether the area of the county is more or less than 14,693 square miles, which is more than 4 standard deviations above the mean of 1069 square miles. Of those counties whose area is more than 14,693 square miles, 10 were from outside Cluster 0, which is about 9.5% of such counties, and none were from Cluster 0. Here, we see that another subset of the non-Cluster 0 counties have incredibly large areas, and in fact, this accounts for all the particularly large counties.

Of the counties along the branch of this split, for which counties have an area smaller than 14,693 square miles, 9 were from outside Cluster 0, which accounts for the remaining 8.5% of such counties, and 2957 were from Cluster 0, which accounts for the remaining 99.5% of these counties.

With these 3 splits of cumulative COVID-19 case count as of August 8, number of COVID-19 testing locations, and county area, we break off almost 87% of non-Cluster 0 counties (97 counties) while only capturing about 0.4% of Cluster 0 counties (14 counties). We also note that the value of each feature split represents some number significantly above the average value of that feature. This suggests that Cluster 0 might represent those counties whose features values are within a reasonable range as seen in the data, and those counties outside Cluster 0 represent outliers with respect to different features. This result could be due to the way that agglomerative clustering operates. Recall that it successively combines the nearest clusters, and so if one cluster starts to represent typical counties, then we would expect it to very often be the nearest cluster to some other county or cluster. This would explain why we see one particularly large cluster and then a few smaller clusters containing mostly outliers.

### 5.5.2 Distinguishing the Smaller Clusters

Now, we want to understand what distinguishes these remaining smaller clusters outside of Cluster 0, including the singleton clusters. We first note that the features with the highest variance among the remaining counties are, in order: (1) population (average outside Cluster 0: 1,121,476), (2) the number of cumulative COVID-19 cases as of August 8, which is the date on which we gathered the data (average outside Cluster 0: 22,595), (3) area in square miles (average outside Cluster 0: 4520), (4) number of new COVID-19 cases on July 23, which is around the date of the second US peak in cases (average outside Cluster 0: 3671), and (5) the number of cumulative COVID-19 deaths as of August 8 (average outside Cluster 0: 909). We will explore these features as we discuss the results that agglomerative clustering has generated.

In the agglomerative clustering process, we noticed that if we were to continue combining clusters, we would very often merge clusters with Cluster 0. Since many of our clusters contain only one or a few counties, in order to group them in a meaningful way for interpretation, we trained a decision tree to distinguish all clusters excluding Cluster 0. This allowed us to see which clusters landed in nearby leaf nodes, suggesting that they represent counties with similar feature values. In Figure 16 we show a simplified version of the resulting decision tree where we have grouped the clusters into 8 groups. We disregard clusters whose counties were split into a number of small groups across our decision tree, and focus only on clusters that remained completely or mostly intact within a leaf node. Across our 8 groups of clusters, we account for 22 of the remaining 24 clusters, which contain 102 of the remaining 106 counties.

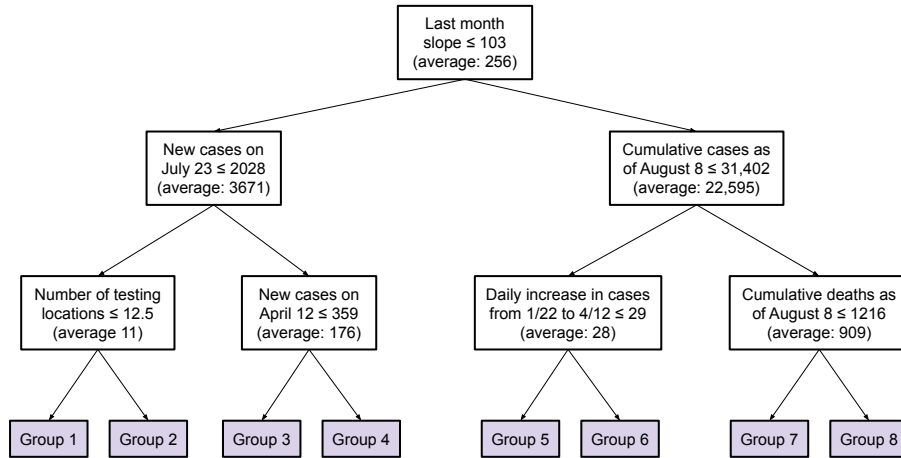


Figure 16: Simplified Decision Tree Distinguishing Groups of the Smaller Clusters.

Group 1 in our simplified decision tree contains 3 clusters, totalling 10 counties. Of these, seven are counties in Alaska, 2 are in Nevada, and one is in Arizona. With respect to the features that vary the most across non-Cluster 0 counties, the counties in Group 1 all have populations well below average; areas well above average; and new cases on July 23, total cases as of August 8, and total deaths as of August 8 well below average.

Group 2 contains 17 of the 18 counties from a single cluster and no other counties. The counties in this group are all in Massachusetts, Vermont, and Connecticut. Comparing the counties in this group to the general counties outside Cluster 0 with respect to the most variable features, we see that Group 2 has values for each (population, cumulative case count as of August 8,

area, new cases on July 23, and cumulative death count as of August 8) that are below average. We split this group apart from Group 1 because the counties in Group 2 have a significantly smaller area than the counties in Group 1, and because for the remaining features, the counties in Group 1 seem to have much more extreme values.

Group 3 contains one cluster of size 5 and 11 counties from a cluster of size 12 for a total of 16 counties. The cluster of size 12 includes counties from across the country that all either include a larger city or are near a larger city, like Orleans, LA (New Orleans) and Macomb, MI (bordering Detroit). The cluster of size 5 contains a few counties in Massachusetts and a few in Connecticut. The counties in this group have populations below average, areas that are well below average, and total cases as of August 8 that are below average.

Group 4 contains a number of very small and singleton clusters, totalling 11 counties. These include 3 of the boroughs of New York City: Bronx, Queens, and Kings (Brooklyn). The other counties are all in New York (outside New York City) and New Jersey. The counties in this group all have well below average area, above average cases on July 23, and above average cumulative number of deaths as of August 8. Recall that New York City was hit really hard with COVID-19 cases at the beginning of the pandemic, and so it makes sense to see it as well as some of the surrounding counties grouped together.

Group 5 contains 28 of the counties from a cluster of size 31. The counties in this group are from all over the country and include Hennepin, MN (Minneapolis), Cuyahoga, OH (Cleveland), Milwaukee, WI, and Alameda, CA (Oakland). We note that the counties in this group have a population barely below average. They also have below average values for area, case count on July 23, total cases as of August 8, and total deaths as of August 8. Overall, this group represents counties with somewhat average feature values, and in fact, in the dendrogram for our hierarchical clustering, this group was fairly close to Cluster 0. This means that within a few more merges of agglomerative clustering, this group would have been combined with Cluster 0, which was the cluster representing counties with less extreme values.

Group 6 contains one singleton cluster. In particular, it contains the county of Wayne, MI (Detroit). We note that Wayne, MI has a population that is somewhat larger than average, and an area that is below average. We separated it from Group 5 because its values for case count on July 23, total cases as of August 8, and total deaths as of August 8 were all above or well above average. While this county may have non-COVID-19 related features with similar values to Group 5, for the features that describe COVID-19 case count and death count, Wayne, MI is an outlier.

Group 7 contains one cluster of size 8 and one singleton cluster, for 9 total counties. The singleton cluster is of San Bernardino, CA, and the other cluster contains counties that all contain a larger city, including Clark, NV (Las Vegas), San Diego, CA, and Dallas, TX. The counties in this group have larger than average populations, fewer than average cases on July 23 and fewer than average total deaths as of August 8, but more than average total cases as of August 8. Due to the fact that these counties have fewer cases in July and fewer deaths at the beginning of August but more cases at the beginning of August, this suggests that these counties may be at the beginning of a spike in cases.

Finally, Group 8 contains one cluster of size 2 and 3 singleton clusters, for 5 total counties. The singleton clusters contain the counties Harris, TX (Houston), Los Angeles, CA, and Cook, IL (Chicago). The cluster of size 2 contains Maricopa, AZ (Phoenix) and Miami-Dade, FL. This group differs from Group 7 in that the counties in Group 8 have significantly higher case counts on July 23 and significantly higher total deaths as of August 8. Additionally, the counties in

Group 8 all have significantly higher than average populations,, and significantly higher total cases as of August 8. These counties represent some especially hard hit counties in the US, particularly over the summer.

### 5.5.3 Summary of Hierarchical Clustering Results

In our case, we see that agglomerative clustering creates one very large cluster that contains most of the counties in our dataset, and the remaining clusters tend to contain outliers with respect to some subset of our features. This could be due to the fact that we have relatively few features compared to the number of counties we have, and the values of these features are somewhat correlated, in particular for the COVID-19-related features. We attempted to mitigate this by applying PCA before clustering.

Changing the final total number of clusters in our agglomerative clustering mostly changed the number of clusters that were added to our one large cluster. Thus, after our initial cut off of 25 clusters, we further group together the remaining clusters using a decision tree. As we saw, these groups can be distinguished from each other based on their values for some subset of the 5 most variable features and by following the decision tree’s splits. The groupings we see in our tree seem fairly reasonable. For example, we see one group that contains most of the counties with the largest land area, another that contains counties in and around New York City, and another that contains many of the counties with the largest populations.

## 5.6 OPTICS Clustering

In this section we describe our work implementing the OPTICS clustering algorithm. We also discuss and interpret the clustering results.

### 5.6.1 Dimension Reduction and OPTICS Clustering

Before clustering our data, we applied the dimension reduction techniques PCA and t-SNE to the dataset. The PCA process is described above in Section 5.5 and t-SNE is similar, although it by definition reduces the dimensions of the dataset to either two or three dimensions. We generated two datasets with two and three attributes, respectively, using t-SNE. We also generated three datasets with three, five, and ten attributes by using PCA which store around 70%, 85% and 97% of the variance in the original data, respectively.

For each of these five datasets, we applied the grid-search function we mentioned in the Section 4.6.2 to find the best set of parameters for the OPTICS clustering algorithm, as well as the best number of clusters to use given those parameters. The results of this grid-search returned higher values for both silhouette score and Calinski-Harabasz index for a large number of clusters for the datasets that we applied t-SNE to but returned higher values for both scores for fewer than five clusters for the datasets that we applied PCA to. Since we expected that ending up with fewer than five clusters would not produce especially useful results, we set conditions on our grid search to account for this and selected the resulting best combination of the values of parameters for each of these datasets, which we show in Table 1.

Based on the silhouette score and Calinski-Harabasz index, as well as the fact that this results in a fairly reasonable number of clusters, we decided to focus on the dataset to which we applied PCA in order to reduce to 5 features. In addition, we chose this dataset since with 5 principal

	Number of dimensions	Min samples	Perplexity	Number clusters	Silhouette score	Calinski-Harabasz score	Model matrix	Silhouette score matrix
t-SNE	2	11	40	132	0.4101	59.5546	cosine	cosine
t-SNE	3	14	5	68	0.0080	57.2414	cosine	cosine
PCA	3	9	/	75	-0.3464	31.5931	cosine	cosine
PCA	5	9	/	32	-0.6436	78.9359	cosine	cosine
PCA	10	19	/	6	-0.2871	269.5029	cosine	cosine

Table 1: The Value of Parameters Selected from the Dataset after Applying Dimension Reduction Techniques

components, we are able to explain 85% of the variance in the original data.

The distribution of counties per cluster for this model is shown in Table 2, and a map with all counties, excluding those HI or AK colored by cluster number is shown Figure 17.

Labels	Number of Counties	Labels	Number of Counties	Labels	Number of Counties
-1	2579	10	40	21	10
0	15	11	14	22	16
1	19	12	19	23	13
2	14	13	16	24	10
3	14	14	9	25	21
4	15	15	16	26	16
5	12	16	14	27	17
6	21	17	13	28	14
7	11	18	13	29	17
8	12	19	19	30	14
9	16	20	21	31	19

Table 2: Distribution of Clusters

We can see the majority of counties are labeled as noise, corresponding to cluster -1, shown in Figure 17. This can be explained by the fact that the most of the confirmed cases and death are concentrated in a small number of counties while the values for these features in most of the other counties are extremely small. This causes large distances between these counties in the dataset. Thus, the counties where these values are small are prevented from being assigned a cluster and are labeled as outliers.

### 5.6.2 Summary of OPTICS Results

After we removed the outliers, which correspond to counties in cluster -1, we wanted to see how the values of the features are similar or different across the remaining clusters. To do this, we calculated the standard deviation for each attribute and removed those with relatively small standard deviations since they represent features with similar values across all clusters. Then for each cluster, we computed the mean value of the 12 attributes in our data with larger standard deviations and plotted these values along with the cluster name in Figure 18.



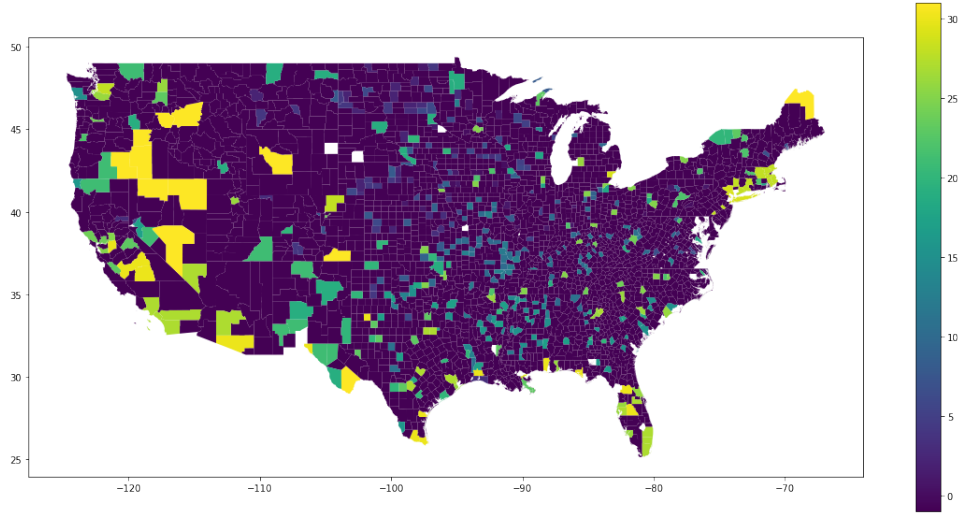


Figure 17: Location of the Counties in Each clusters

We can see from Figure 18 that for clusters 0 through 23, as well as for cluster 31, the values of most attributes are fairly constant and also are slightly negative. On the other hand, clusters 24 to 30 have extremely large value of each attributes compared with other clusters. Of the 12 attributes that we have plotted in Figure 18, 11 have their highest average value in cluster 27, cluster 28 and cluster 29. In the discussion below, we will analyze these three clusters to figure out what caused their unusual values.

Cluster 27 contains 17 counties located in 6 different states. We have listed these counties in Table 3 and shown them on a map in Figure 4. Most of the counties in this cluster are concentrated at California, Texas and Florida and are located in the southern part of the US. This cluster has the highest mean value of area, population, number of ICU beds and number of cumulative cases as of August 8. Also, it has an extremely high mean value of the case growth rate for the month leading up to August 8, compared with the rest of the clusters. Additionally, the majority of these counties have high populations compared to other counties in their state. For example, Harris, Dallas, Tarrant, Bexar and Travis are the top 5 counties in population in Texas. We know that the states of California, Texas, and Florida all have high numbers of confirmed cases after June 2020. With the above information, we could conclude that this cluster contains the counties with most rapid growth of the number of confirmed cases and deaths after June.

Cluster 28 contains 14 counties located in 7 different states. We have listed these counties in Table 5 and shown them on a map in Figure 6. We can see that most of these counties are in Massachusetts, Connecticut, New York, New Jersey, and Pennsylvania. Almost all the counties are located in the northeast part of the US except King, WA and Weld, CO. This cluster has an average population of 799,174, has the highest average number of testing locations, and has a relatively high nursing home population. This indicates the counties in this cluster have good medical resources and focused on stopping the spread of COVID-19 by adding more testing locations. This cluster also contain counties with large populations and relatively low cumulative confirmed cases as of August 8 compared with other clusters.

Cluster 29 contains 17 counties located in 4 different states. We have listed these counties in Table 7 and shown them on a map in Figure 19. Fifteen of these counties are in New York and

Cluster	Location
27	Hillsborough, FL
27	Palm Beach, FL
27	Maricopa, AZ
27	San Diego, CA
27	Broward, FL
27	Orange, FL
27	Travis, TX
27	Los Angeles, CA
27	Orange, CA
27	Riverside, CA
27	Gwinnett, GA
27	Clark, NV
27	Tarrant, TX
27	Miami-Dade, FL
27	Bexar, TX
27	Harris, TX
27	Dallas, TX

Table 3: County List of Cluster 27

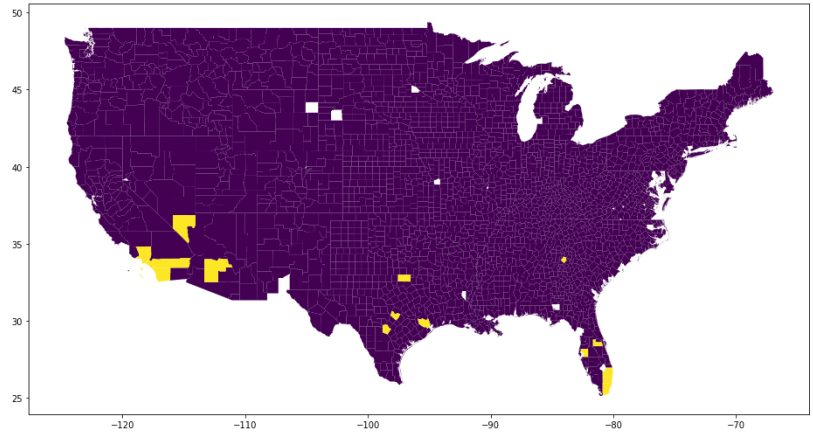


Table 4: County Location of Cluster 27

Cluster	Location
28	Plymouth, MA
28	Worcester, MA
28	Dutchess, NY
28	Weld, CO
28	Bristol, MA
28	Norfolk, MA
28	Fairfield, CT
28	Hartford, CT
28	New Haven, CT
28	Middlesex, MA
28	Delaware, PA
28	King, WA
28	Essex, MA
28	Mercer, NJ

Table 5: County List of Cluster 28

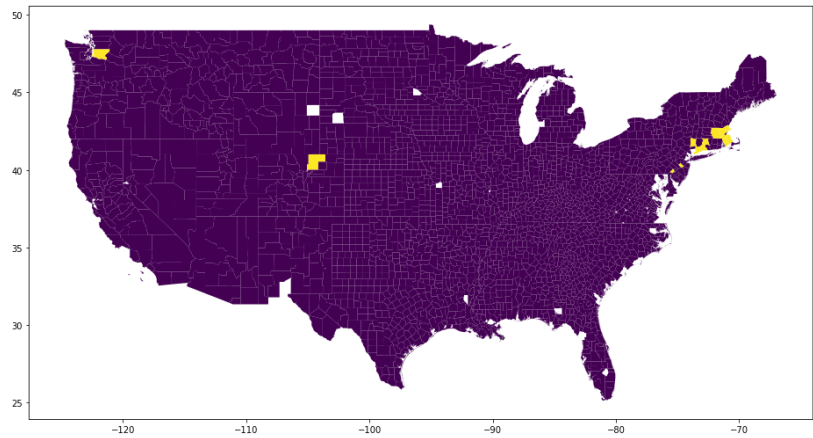


Table 6: County Location of Cluster 28

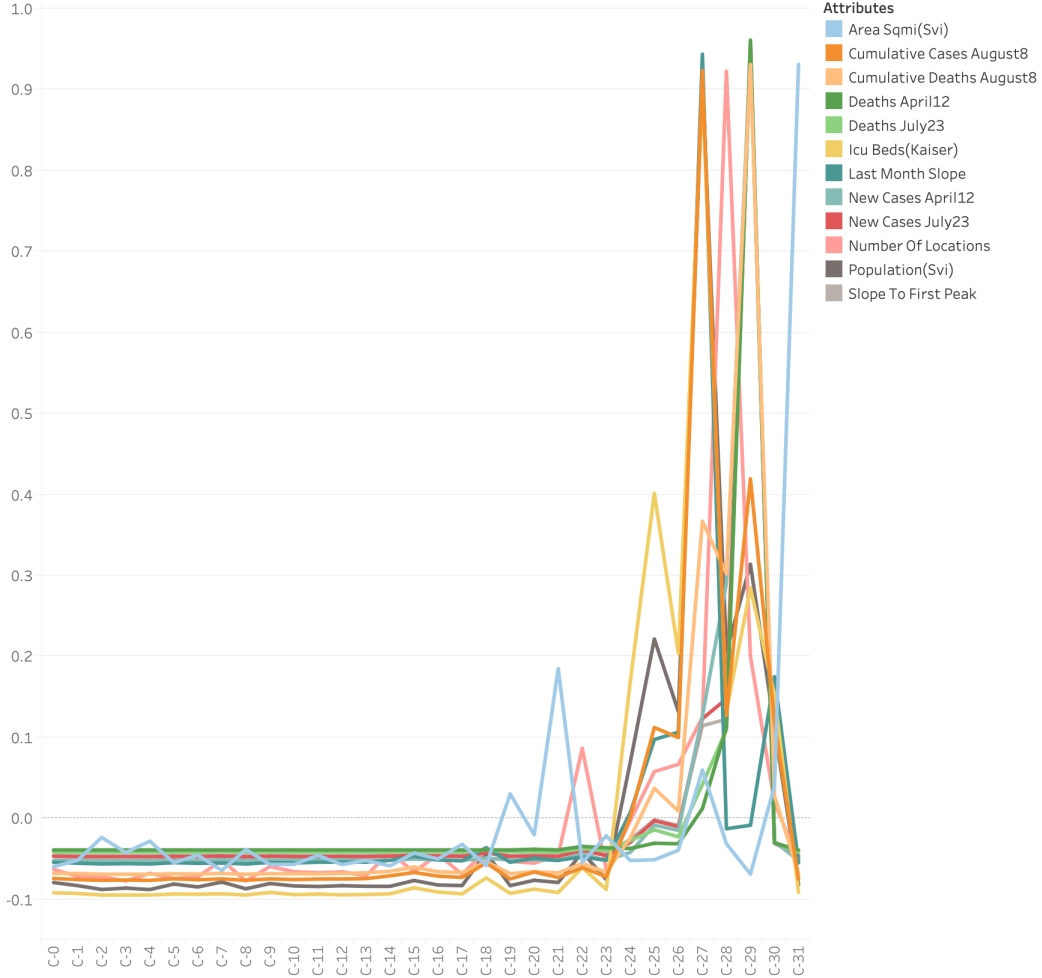


Figure 18: The Mean Value of Each Attribute with Clusters

New Jersey. This cluster has the highest average number of new cases on April 12, the highest average number of cumulative deaths as of August 8, and the highest average growth rate for the first month of the pandemic. Based on this information, we see that this cluster grouped the counties with the highest number of confirmed cases and deaths before June. This agrees with the fact that when COVID-19 first started to spread in the US, New York and New Jersey had particularly high numbers of cases.

## 6 Conclusion

Among the prototype-based methods, we saw that  $K$ -means and Fuzzy- $c$  means clustering methods produced somewhat similar results, which makes sense since they are similar algorithms. Both approaches clustered many counties in the southern half of the country with counties in the west into one cluster, counties in the northern half of the country into another cluster, and some seemingly outlier counties into another few clusters. The Gaussian mixture model and Mini Batch  $K$ -means produced similar clusters to each other. Both had one very large cluster of counties and another cluster or two containing far fewer counties.

Cluster	Location
29	Suffolk, NY
29	Nassau, NY
29	Orange, NY
29	Bergen, NJ
29	Richmond, NY
29	Westchester, NY
29	Orleans, LA
29	Wayne, MI
29	New York, NY
29	Hudson, NJ
29	Passaic, NJ
29	Union, NJ
29	Kings, NY
29	Queens, NY
29	Rockland, NY
29	Essex, NJ
29	Bronx, NY

Table 7: County List of Cluster 29

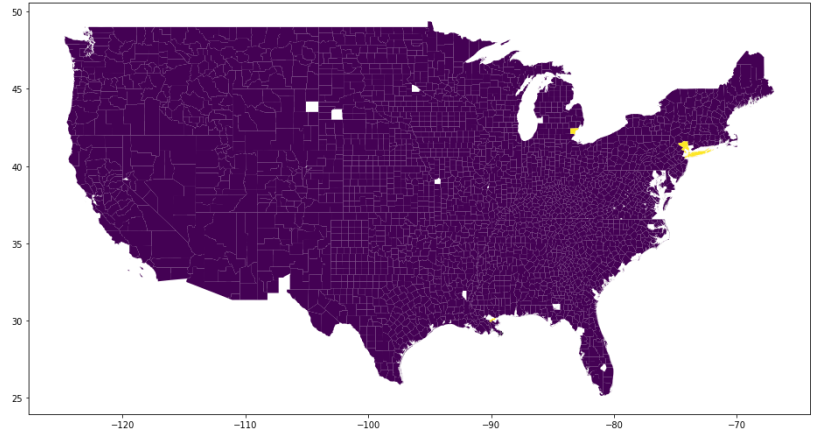


Figure 19: County Location of Cluster 29

Across these four prototype-based methods, we saw almost all of our features were roughly equally as useful in distinguishing our clusters, suggesting that all of our features could provide useful insight into our data. According to the important features selected by each of the prototype-based approaches, we saw that a county’s population, index of relative rurality, number of ICU beds, and ranking with respect to socioeconomic status came up most frequently. Thus, we conclude that these features are closely related to how severely the COVID-19 pandemic impacted those counties. The features describing the county ranking with respect to housing and transportation as well as the county ranking with respect to household decomposition and disability also seem related to the severity of the COVID-19 pandemic in a given county to some extent.

The hierarchical clustering and the density-based clustering via OPTICS both output one very large cluster of counties that we can view as “noise” or counties with more average values for each of our features, and a number of much smaller clusters. In hierarchical clustering, we found a number of singleton clusters, and only a few clusters with sizes in between. Due to the nature of the OPTICS algorithm, the smaller clusters all contained at least a pre-specified number of counties, resulting in more uniform sizes for these smaller clusters. Since we were ultimately interested in finding similar counties, this suggests that perhaps the OPTICS algorithm is better suited between these two methods.

Overall, from this project, we found that applying clustering to US county data related to the COVID-19 pandemic revealed a variety of insights. Due to the nature of clustering, we did not select one best approach, and instead found that each technique revealed different patterns and structure within our data.

## References

- [1] Mihael Ankerst et al. “Optics”. In: *Proceedings of the 1999 ACM SIGMOD international conference on Management of data - SIGMOD 99* (1999). DOI: 10.1145/304182.304187. URL: <https://www.dbs.ifi.lmu.de/Publicationen/Papers/OPTICS.pdf>.
- [2] David Arthur and Sergei Vassilvitskii. “K-Means++: The Advantages of Careful Seeding”. In: *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms. SODA '07*. New Orleans, Louisiana: Society for Industrial and Applied Mathematics, 2007, pp. 1027–1035. ISBN: 9780898716245.
- [3] Thomas Charlon. *opticskxi: OPTICS K-Xi Density-Based Clustering*. URL: <https://cran.r-project.org/web/packages/opticskxi/vignettes/opticskxi.pdf>.
- [4] *DBSCAN*. URL: <https://en.wikipedia.org/wiki/DBSCAN>.
- [5] *Sci-Kit Learn User Guide*. URL: [https://scikit-learn.org/0.21/user\\_guide.html](https://scikit-learn.org/0.21/user_guide.html).
- [6] Pang-Ning Tan et al. *Introduction to data mining*. Pearson Education, Inc., 2019.