

Measuring Airport Similarity to Create a Towering Decision Aid

A Technical Report submitted to the Department of Engineering Systems and Environment

Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

Michael Quinn
Spring, 2020

Technical Project Team Members

Austin Anderson
Toby Hansford
Mason Jordan
Sragi Khakurel
Chris Marshall
Katherine Taylor
Amy Xie

On my honor as a University Student, I have neither given nor received
unauthorized aid on this assignment as defined by the Honor Guidelines
for Thesis-Related Assignments

Signature Michael Quinn Date 05/01/2020
Michael Quinn

Approved _____ Date _____
Cody Fleming, Department of Engineering Systems and Environment

ABSTRACT

The focus of this project was on formulating a model and decision support tool to aid in the decision to build and maintain an Air Traffic Control Tower (ATCT). An important aspect of air travel are ATCTs, towers that help facilitate communication between the airport system and airplanes ascending and descending. ATCTs bring economic, safety, and efficiency benefits to airports and nearby communities. Currently, the Federal Aviation Administration (FAA) uses a document outlining a benefit-cost ratio for building a new tower, with tower funding provided if the ratio is greater than 1. However, the current policy lacks a comprehensive and systematic assessment of factors that influence both costs and benefits to operators and the region.

To address these issues, we started by speaking with air traffic stakeholders and then began to collect data from a variety of aviation datasets. Based on the collected data, we identified economy, safety, and efficiency as our three areas of focus. With this data, we were able to compute the similarity, using hierarchical clustering, of a given airport to currently towered airports based on data from the economy, safety, and efficiency sources. We then built an interactive interface to display these similarities and provide information for airports to contact the similar airports.

INTRODUCTION

Air Traffic Control Towers (ATCTs) have the primary responsibility for preventing collisions between aircraft and other hazards (*Air Traffic Control*, n.d.). In the United States, there are approximately 500 towered airports, and 20,000 non-towered airports (S.D. Van Beek, July 2019). Of these approximately 500 towered airports, only 264 are directly run by the Federal Aviation Administration (FAA), the rest are contracted out, at the cost of 26 percent of an FAA tower. To establish a new tower, the FAA has several criteria, the most important being

a benefit-cost analysis of the tower establishment. However, the FAA has not rerun the benefit-cost ratios for existing towers since 2006, and its methodology is biased towards airports with higher operational volumes.

The benefit-cost analysis method the FAA uses to determine whether an airport needs a tower is outdated and inflexible. As the FAA has constrained resources and budget, and the cost to build a physical tower increases, smaller airports may desire more information about the benefits a tower would bring to their airport. We propose an interface that helps smaller airports interested in towering find similar airports based on metrics about economics, safety, and efficiency. To determine which airports are most similar, we used a hierarchical clustering algorithm.

BACKGROUND

Air travel has become increasingly important in the United States for both business and leisure. In turn, issues with congestion, funding, and the environment are more prevalent than ever before. Therefore, updated air traffic control schemes and infrastructure are needed to improve the shortcomings of the current system. ATCTs are a service provided at airports to help improve and control air traffic through direction and advisory services (*Air Traffic Control*, n.d.). ATCTs have a clear benefit of preventing collisions and allowing for more efficient flights at larger airports with commercial carriers. The FAA builds ATCTs in Class B, C, and D airspaces and generally does not build ATCTs in other airspace classes. Currently, the FAA determines whether or not to fund the building and operating of these towers at airports based on the document Establishment and Discontinuance Criteria for Air Traffic Control Towers (FAA-APO-90-7).

The Establishment and Discontinuance Criteria for Air Traffic Control Towers document outlines a benefit-cost ratio, comprising safety and efficiency as benefit factors (Martin, G.A., n.d.). If the ratio is above 1, the tower will be funded, if below 1; the tower will not be funded. The existing guidelines put an emphasis on the benefits outweighing the costs. This constitutes a problem when both the benefits and the costs are very high or small. The ratio is also outdated in terms of benefits and costs, as it was last updated in 1990. Additionally, because the current criteria are outdated, they are biased against the class D airspace. This may make it difficult for small airports to obtain funding for a tower. However, smaller airports are in need of these towers as they not only lead to increased safety and efficiency, but they also economically benefit the surrounding community from additional commercial and corporate traffic.

APPROACH

A. Data Collection

Based on conversations with our client and other professionals in the aviation industry, our team decided to focus on 3 metrics: efficiency, safety, and economy. All of our data was collected from online databases, and include data through December of 2019. The efficiency data, consisting of volume data for different types of flights, was from the FAA's Operations System Network (OPSNET) database and included all flights in the calendar year 2019. The safety data was from a series of incident reports recorded by the National Transportation Safety Board (NTSB) dating back to 1982. The information contained in each report includes the date, location, fatalities, and aircraft damage. We consolidated the safety data to include one entry per airport by aggregating the incident reports. From data provided by the client, we labeled which airports had ATCTs, and removed accidents prior to the towering date. Lastly, the economic data was collected from various state aviation economic impact studies and includes data from 2003

to 2018. We were able to find economic output, employment, and wage data on airports in 26 of the 50 states. The data was then adjusted for inflation to be in terms of 2019 dollars. We compiled all the data into a data frame for each of our metrics using Microsoft SQL server and created a combined dataset based on airport code. We were able to find data on all of our metrics for 228 airports.

B. Analysis Method

We chose to limit our analysis options to unsupervised learning techniques as we did not have a clear feature label for our airports. Unsupervised techniques allowed us to study the similarities between airports in terms of our metrics. Clustering algorithms lend themselves well to unsupervised learning, as they look more at the similarities and dissimilarities of the data rather than a fixed output. Cluster validation showed hierarchical clustering to be the best method for performing unsupervised learning on our data, over k-means clustering. Hierarchical clustering has the additional bonus of not requiring the number of clusters to be specified prior to running the algorithm, and the result can be split into clusters based on visual inspection. To conduct our analysis, we followed the method in the textbook *Practical Guide to Cluster Analysis in R* (A. Kassambara, 2017).

Hierarchical clustering is a “recursive partitioning of a dataset into successively smaller clusters,” (V. Cohen-Addad et al., 2019). The method of hierarchical clustering can be split into two main phases, similarity analysis, and tree construction (P. Daie & S. Li, 2016). The input into the algorithm consists of a matrix with pairwise similarities or dissimilarities between the airports. For our analysis, we used Euclidean distance to compute pairwise dissimilarities between the airports. Hierarchical clustering tree construction algorithms generally fall into one of two categories, agglomerative or divisive. Agglomerative methods build the tree starting at the

node level, while divisive methods split the tree starting from the top (V. Cohen-Addad et al., 2019). For this analysis, we considered only the agglomerative methods of complete-linkage, Ward's linkage, and average-linkage. For an agglomerative method, each observation starts out as a single cluster. The two clusters that are the closest to each other using a distance metric are then joined by the distance metric. This step is repeated until all observations are included in the same cluster. Linkage functions define how clusters should be joined past the first cluster containing two elements. Complete linkage uses the maximum distance from the cluster to other observations, while average linkage uses the average distance. Ward's linkage method attempts to form clusters by minimizing the within-cluster distances while maximizing the between-cluster distances. Ward's method joins two clusters A and B that minimize the increase in the sum of squared errors (SSE).

$$I_{AB} = SSE_{AB} - (SSE_A + SSE_B) \quad (1)$$

SSE_{AB} is defined as the between-cluster sum of squared errors, while SSE_A and SSE_B are the within-cluster sum of squared errors. This objective function can also be written as

$$I_{AB} = \frac{n_A n_B}{n_A + n_B} (\bar{a} - \bar{b})'(\bar{a} - \bar{b}) \quad (2)$$

Where \bar{a} and \bar{b} are the centroids of clusters a and b (T. Strauss & M. Johan von Maltitz, 2017).

The results of using a hierarchical clustering algorithm can be visualized using a dendrogram, which displays the distances between each of the observations, and from which specific clusters can be identified. An example figure of a dendrogram is given below in Fig 1.




Figure 1 Example Dendrogram with Colored Clusters (Data: USArrests)

This sample dendrogram was created using the USArrests dataset, which includes the number of arrests for various violent crimes by state. In the above dendrogram, four different clusters are colored, and the states that are a part of each cluster are labeled. States that are connected at a smaller height value are more similar, and states that are connected larger are less similar. We created a dendrogram for this analysis in order to better visualize the similarities between the airports.

DATA ANALYSIS & RESULTS

A. Analysis

The dataset we used for the analysis contained the columns: jobs, wages, economic output, number of accidents, deaths, serious injuries, minor injuries, uninjured, number of years in the aggregation, years since last accident, and operation counts for air carrier, taxi, general aviation, itinerant military, civil, and local military flights. The commercial status of the airport

was not used in performing the principal components analysis (PCA) or clustering because it is a categorical variable. All of the data used in the PCA and clustering analysis was scaled.

The principal components analysis plot based on the above columns is shown below in Fig. 2. The majority of airports are clustered in the left of the PCA plot, with very few having large values of component one. Looking at Table 1, this means that the outliers have significantly more jobs, wages, economic output, number of accidents, deaths, serious injuries, minor injuries, uninjured, number of years in the aggregation, and operation counts for air carrier and taxi flights. They have significantly less years since the last accident, and operation counts for general aviation, itinerant military, civil, and local military flights. That being said, to improve the generality of the analysis, airports with an outlying first principal component were excluded.




Fig. 2. Principal Components Analysis

In order to use a clustering algorithm, it must first be shown that there are significant clusters in the data either through the Hopkins test or the visual assessment of cluster tendency (VAT). The Hopkins test is a statistical hypothesis test where the null hypothesis states that the data are uniformly distributed. The alternative hypothesis states that there are significant clusters in the data. The threshold for rejecting the null hypothesis using the test statistic H is 0.5, values that are closer to 1 correspond to data sets that are uniformly distributed, and values that are closer to 0 contain meaningful clusters. Since the data has a Hopkins test statistic of 0.0533, which is less than the threshold of 0.5, the data is considered to have significant clusters. The result of the Hopkins test is confirmed using the VAT, as seen through the patchwork pattern in Fig. 3. The clearly defined red squares indicate there are significant clusters in the data. If the majority of the patchwork was blue, then there would not be significant clusters in the data.




Fig. 3. Visual Assessment of Cluster Tendency

After performing a comparison of k-means clustering and hierarchical clustering, hierarchical clustering with two clusters was found to be the best fit for the data. The results of the comparison using three metrics are shown below in Table 2.

	Score <dbl>	Method <fctr>	Clusters <fctr>
Connectivity	8.1544	hierarchical	2
Dunn	0.3305	hierarchical	2
Silhouette	0.5396	hierarchical	2

Table 2: Optimal Clustering Method

The hierarchical clustering was carried out using Ward’s D linkage, which is described in the Approach section. The graphical visualizations of the two resulting clusters are below with an abstracted dendrogram in Fig. 4, and a cluster plot in Fig. 5. Additionally, the cluster means are displayed in Table 3. The main difference between the two clusters is that the second cluster has significantly higher values of every variable except general aviation and civil operations. An expanded version of the second cluster is shown below in Fig. 6.

Cluster Dendrogram

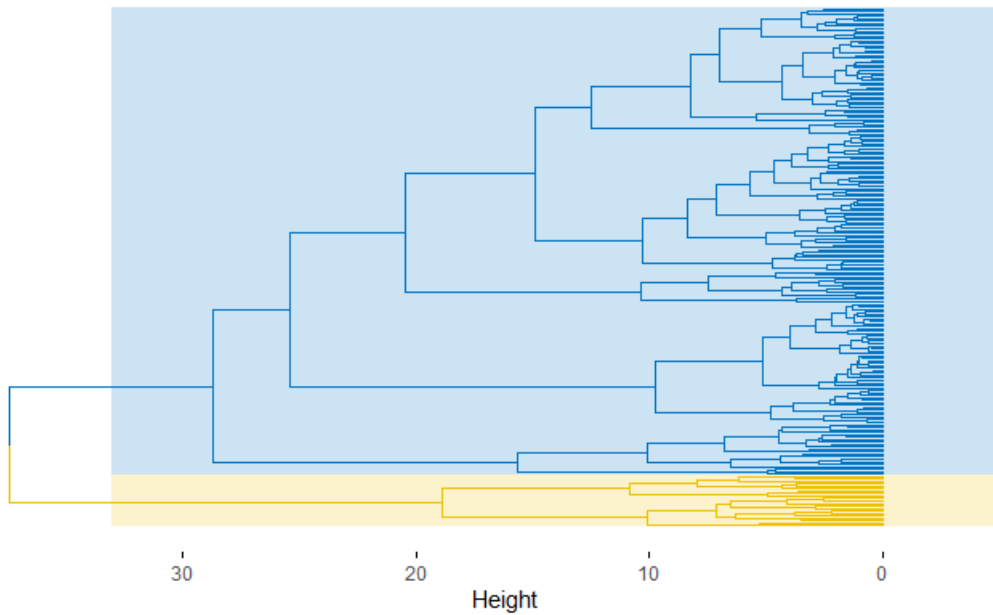


Fig. 4: Full Dendrogram




Fig. 5: Cluster Plot Derived from the Hierarchical Clustering

cluster	jobs	inflation_wages	inflation_output	number_accidents	total_deaths	total_serious	total_minor	total_uninjured
1	3063.753	132392076	425820783	16.71579	5.384211	2.305263	3.994737	46.41053
2	38358.381	1484013340	5050143242	24.42857	11.000000	4.000000	8.238095	380.52381

number_of_years	year_since_accident	air_carrier	taxi	general_aviation	itinerant_military	civil	local_military
28.29474	4.289474	3913.832	9846.442	29649.52	2789.584	27217.300	2268.032
36.76190	5.142857	62734.667	17571.571	30696.67	4657.048	9489.429	2083.143

Table 3: Cluster Means for each Variable




Fig. 6: Dendrogram of Cluster 2

It can be seen from the full dendrogram that the airports in the second cluster are more similar to each other in the first cluster, and there are only 21 airports in the second cluster while there are 190 airports in the first. Within the second cluster, there are two distinct subclusters, one which appears to have airports serving larger regions such as Houston (HOU) or New Orleans (MSY), and the other subcluster contains airports serving smaller cities such as Jacksonville (JAX), or mixed civilian and military airports like Savannah and Hilton Head Airport (SAV). From the cluster plot (Fig. 6), there is some overlap between the two clusters, which means that some elements from the second cluster could belong in the first cluster.

The silhouette plot/index can be used to internally validate the groupings for each of the clusters. A negative silhouette index indicates that the observation potentially belongs in a different cluster. The silhouette plot is below in Fig. 7. From this plot, it can be seen that a few members of the first cluster do not belong in the cluster, and approximately a fourth of the members in the second cluster do not belong. Overall, the hierarchical clustering performs

relatively well with an average silhouette distance of 0.43, indicated by the red dashed line on the plot.




Fig. 7: Silhouette Method Plot

The clustering analysis described above was then used to design an interface to connect airports looking to build a new ATCT to similar airports with existing ATCTs.

B. Results/Deliverables

With the clustering results from above, we decided to build a user interface system that would easily allow key stakeholders to access this information, and allow for the individual adjustment of desired metrics. The interface gives users background statistics on any airport they select as well as a visualization for the similarities among different airports. We decided to build this interface using Tableau so that the data can be easily visualized. In addition, Tableau can connect to R, which provides the ability to bring statistical analysis into a visual analytics environment.

The interface itself is divided into two sides, left and right (Fig. 8). By selecting an airport code from the dropdown menu on the left-hand side, the interface returns various

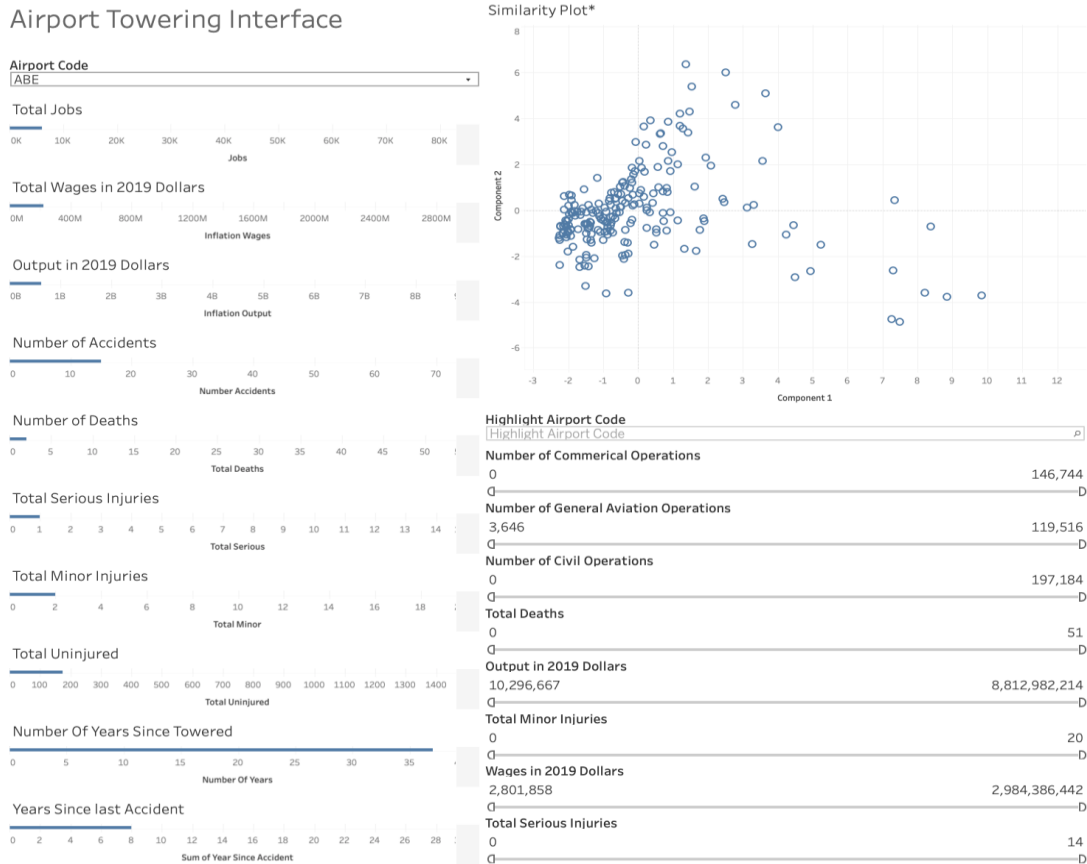


Figure 8 Interface Design

economic, efficiency, and safety statistics on the airport in the form of bar charts. There is a total of sixteen different metrics, including the number of jobs the airport employs, total wages it outputs, number of accidents it has experienced, number of operations it carries out annually, and more. On the right-hand side, there is a similarity plot as well as the same sixteen statistics, but in the form of slider bars. The user can adjust these sliders to narrow down the data points on the similarity plot, showing the user which airports fall into the specified range of statistics as well as how similar the displayed airports are to each other. Furthermore, the user can select a

particular airport code from a dropdown menu below the similarity plot to highlight that airport's exact position on the similarity plot.

The visualization of the data for each airport, as well as the ability to adjust metrics to resemble an airport in question, aids the user in deciding whether a tower is appropriate or not for said airport. As opposed to simply returning a number or ratio, this provides a more comprehensive and in-depth view into the intricacies of the decision to tower.

CONCLUSION

In performing PCA and hierarchical clustering analysis, we produced an interface that aids in the decision of whether an airport should establish a tower. The clustering analysis groups airports using a given set of variables, and with this technique, we are able to provide users with a list of comparable airports that have towers, as well as information about the airports' economic, safety, and efficiency; the interface also provides users with a point of contact for comparable airports. These resources provide the user with a greater insight into potential benefits associated with having an ATCT. It also addresses the robustness issues associated with the current establishment criteria through the capability of dynamic data and ease of statistical analysis, which allows the interface to be updated with new data as it is collected. As the FAA has constrained resources and budget, and the cost to build a physical tower increases, this updated decision aid is integral in providing more information to airports looking to build an ATCT.

REFERENCES

- Federal Aviation Administration, *Air Traffic Control*. Washington DC: Air Traffic Organization, 2019.
- S. D. Van Beek “Remote towers: A better future for America’s small airports,” *reason.org* para. July 2019. [Online]. Available: https://reason.org/wp-content/uploads/files/air_traffic_control_remote_towers.pdf. Accessed Mar. 26, 2020].
- Martin, G. A. (n.d.). Faa Contract Tower (Fct) New Start and Replacement Tower Process. FAA Contract Tower (FCT) New Start and Replacement Tower Process.
- A. Kassambara, *Practical Guide to Cluster Analysis in R*, STHDA, 2017.
- V. Cohen-Addad, V. Kanade, F. Mallmann-Trenn, and C. Mathieu, “Hierarchical clustering: Objective functions and algorithms,” *Journal of the ACM*, vol. 66, no. 4, pp. 261-303, 2019.
- P. Daie and S. Li, “Managing product variety through configuration of pre-assembled vanilla boxes using hierarchical clustering,” *International Journal of Production Research*, vol. 54, no. 18, pp. 5468-5479, 2016.
- T. Strauss and M. Johan von Maltitz, “Generalising Ward’s method for use with Manhattan distances,” *PLoS ONE*, vol. 12, no. 1, pp. 1-21, 2017.