



Doorway to the United States: An Exploration of Customs and Border Protection Data

Philippe Monmousseau, Aude Marzuoli, Christabelle Bosson, Eric Féron,
Daniel Delahaye

► To cite this version:

Philippe Monmousseau, Aude Marzuoli, Christabelle Bosson, Eric Féron, Daniel Delahaye. Doorway to the United States: An Exploration of Customs and Border Protection Data. 38th Digital Avionics Systems Conference, Sep 2019, San Diego, United States. hal-02386102

HAL Id: hal-02386102

<https://enac.hal.science/hal-02386102>

Submitted on 29 Nov 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Doorway to the United States: An Exploration of Customs and Border Protection Data

Philippe Monmousseau^{*†}, Aude Marzuoli[†], Christabelle Bosson[‡], Eric Feron[†], Daniel Delahaye^{*}

^{*} ENAC, Toulouse University, Toulouse, France

Email: philippe.monmousseau@enac.fr, delahaye@recherche.enac.fr

[†] School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, USA

Email: amarzuoli3@gatech.edu, eric.feron@aerospace.gatech.edu

[‡] Purdue University, West Lafayette, IN, USA

Email: cbosson@alumni.purdue.edu

Abstract—This paper presents a data-driven study of wait time patterns for international arriving passengers across all 61 terminals from the 44 airports of entry of the United States. Each airport is an independent entity which operates with various airlines and handles demand volumes differently. This induces seasonal variation in service quality from one airport to another. Exploring six years worth of data, this paper investigates the current and long-term performance trends - an increasing number of flights versus a decreasing number of customs booths - of all airports of entry from a passenger perspective. A performance analysis is then conducted that compares average wait times of incoming passengers, considering incoming traffic ratios and allocated resources. Leveraging machine learning algorithms, six regression algorithms are trained and tested to accurately predict passenger wait times through customs at selected airports. An analysis of the performance of these models shows that the best approach - using a Gradient Boosting regressor for each terminal of entry - can capture the daily and seasonal variations of traffic patterns and immigration booth availabilities with a mean absolute error of less or equal to 5 minutes for twenty-eight terminals of entry and less than 10 minutes for all terminals. Observations show significant disparities across airports that may be explained by the foreign/US passenger ratio and the quality of booth management.

Index Terms—big data, passenger-centric, wait time prediction

I. INTRODUCTION

The U.S. Department of Transportation's Bureau of Transportation Statistics (BTS) reported that U.S. airlines carried an all-time high number of passengers in 2016 – 823.0 million systemwide, 719.0 million domestic and 103.9 million international [1]. Passengers are at the core of the Air Transportation System and, yet, obtaining quantitative information about passenger movements remains difficult, particularly on the ground. NextGen [2] in the United States and SESAR [3] in Europe have been advocating a shift from flight-centric metrics to passenger-centric metrics to evaluate the performance of the Air Transportation System. Both NextGen and SESAR intend to not only improve the predictability and resilience of the Air Transportation System, but also to reduce door-to-door travel time for passengers. While passengers whereabouts are known once they board an airplane, their behavior inside the

airport is less understood, and passenger-related information is siloed between different stakeholders, such as airlines, airport security, customs and border protection. Foreign international air travelers arriving in the U.S. spend billions of dollars while visiting. Torres et al. showed that consumption of goods and services grows with the time spent by passengers in the leisure areas. [4]. Thus less time spent queueing at various control points may result in time and financial benefits for everyone.

Airports are known to be the main bottlenecks of the Air Transportation System, yet passenger satisfaction is largely affected by their experience at the airport [5]. Airport management is distributed across various stakeholders, including airlines, federal agents, local airport authorities and third parties, whose decisions may, unbeknownst to them, impact each other's performance [6]. Wait times at security and border patrol play a big role in assessing passenger satisfaction : capacity constraints and inefficiencies at airport entry roads, parking, security, immigration, customs, gates, ramp areas, runways are the primary causes of congestion and of the ensuing delays. Since September 11, 2001, airport screening procedures in the U.S. have been continuously evolving. For example, the passenger screening process is now trying to strike a balance between security and customer service (i.e. minimizing wait times). Using data from 2002 and 2003, Gkritza et al. [7] showed that, while wait times at security screening points are significant determinants of passenger satisfaction, many other factors come into play.

Moreover, delays on the ground have a disproportionate impact for passengers who often experience lengthy delays before being re-booked if they miss their connecting flights [8]. For international passengers arriving from overseas to the United States, immigration checks are mandatory, whether they are U.S. citizens or foreigners.

This paper focuses on the passenger experience while going through U.S. customs and border protection. Roberts et al. studied the evolution of wait times at airports over a few years [9]. They showed that average passport inspection wait time at 24 U.S. airports rose by 25% during 2010-2013. They focused on JFK airport. At JFK, nearly 3 million passengers (25 percent of total arrivals) experienced a delay of more

than 1 hour, putting them at risk of missing a connecting flight, with the 11 percent who had a total delay of more than 2 hours, possibly missing connections at a higher rate. Extended passport inspection waits were the sole source of missed connection risk for 13 percent of passengers and one of the reasons behind missed connections for many more passengers.

There has been little research in the systematic analysis of passenger wait times at customs across airports. Sankaranarayanan et al. [10] performed an exploratory analysis of airport wait times on customs, border protection data taken from top 3 busiest airports (Atlanta, Chicago, and Los Angeles) from the United States, highlighting the effects of seasonality. Johnstone et al. presented a dynamic queue controller to generate realistic queue formation and behaviour within a discrete event environment at airports in Australia [11]. Robertson et al. [12] modeled passenger arrivals to estimate how many passengers arrive at the airport per day and time of day. They also estimated passenger arrival patterns at different processing points (check-in, baggage security, security checkpoint, etc). Nikoue et al. [6] focused on Sydney airport, where they studied flight information, cell phone geolocations from MAC device time stamps, and immigration records from the Australian Department of Immigration and Indigenous Affairs. They modeled arrivals of passengers at Sydney International airport, simulated walk times to immigration and service at immigration queues, to study the effects of a variation in staffing levels at immigration on congestion. However, to the best of the authors' knowledge, no work is available on comparing performance across airports or predicting performance at any airport.

The present paper leverages publicly available data from the United States Customs and Border Protection (CBP) [13]. As stated on their website, "CBP closely monitors the flight processing times, commonly referred to as wait times, for arriving flights at the busiest international airports." The data provided in the online reports show the number of passengers processed on flights arriving in each hour based on how long it took for those passengers to clear Passport Control.

This paper tackles the following research questions:

- Which are the best airports to enter the U.S. for U.S. citizens and foreigners, in terms of wait times?
- Which airports best manage their customs area?
- Can wait times per hour at airports be reliably predicted from historical data?

The paper is organized as follows. Section II explores the main trends at different scales visible in the CBP data, from average wait times, to passenger volumes and flight volumes. Section III compares the performances of the different airports of arrival. Section IV proposes a machine learning approach to predict passenger wait time per hour at any airport and examines the performance of this approach. Section V details the conclusion of the paper and offers future research perspectives.

II. EXPLORATION

A. Dataset contents

The data from CBP [13] contains the following fields once a time period of interest is selected:

- Airport Name,
- Terminal number,
- Date,
- Hour,
- Average wait time for U.S. citizens,
- Average wait time for non-U.S. citizens,
- Maximum wait time for U.S. citizens,
- Maximum wait time for non-U.S. citizens,
- Average wait time for all passengers,
- Maximum wait time for all passengers,
- Number of passengers who waited less than 15 minutes,
- Number of passengers who waited 16 to 30 minutes,
- Number of passengers who waited 31 to 45 minutes,
- Number of passengers who waited 46 to 60 minutes,
- Number of passengers who waited 61 to 90 minutes,
- Number of passengers who waited 91 to 120 minutes,
- Number of passengers who waited over 120 minutes,
- Total number of passengers, both U.S. citizens and non-U.S. citizens,
- Number of flights,
- Number of open immigration booths.

Considering the data from 2013 to 2019, the dataset consists of 1,201,181 entries corresponding to 61 terminals within 44 airports.

B. Long term evolution

Looking at the overall evolution from 2013 to 2019, some clear trends appear as illustrated in Figures 1 & 2. Figure 1 shows the evolution of the total number of arriving international flights per day versus the total number of open booths per day. While the number of flights is steadily increasing over the years, the number of open booths is slowly decreasing.




Figure 1: Comparison of the total number of open booths (red) vs. the total number of arriving flights (blue) per day from January 2013 to January 2019

Figure 2 depicts the daily average wait time per passenger across airports. Wide variations can be noted, from a minimum of 11 minutes to a maximum of 26 minutes. Seasonal

variations are present during the winter and summer holiday season. However, starting from 2015, the amplitudes of these yearly seasonal variations do not vary much over the years.




Figure 2: Average wait time for all passengers per day across all airports from January 2013 to January 2019.

This observation is better visualized in Figure 3, which shows a yearly comparison of the average hourly wait time between 2013 and 2019. With the exception of 2013, this overall average wait time follows the same seasonal variations from one year to another. Longer wait times are observed for the winter (end of December - beginning of January) and summer (August) holidays as well as around April.




Figure 3: Yearly comparison of the average wait time for all passengers per day across all airports

Figure 4 shows the box-and-whisker plot variations of the average wait time across all terminals per month, i.e. for each month it shows the median average wait time and the first and third quartile along with whiskers for a better visualization of the range of the data. This figure highlights the four-month periodic variation discovered in Figure 3. There are three highs during a year - around the winter holidays from December to January, April and August - interlaced with periods with shorter wait times. This four-month periodic behavior does not seem to be induced by a similar behavior in the number of arriving passengers, as shown in Figure 5.

Figure 6 shows the evolution of average wait time for all passengers per hour of the day from 2013 to 2019. As previously observed, the average time for 2013 is higher than all the others, which are tightly packed. The average wait time for 2019 is lower than the previous years since only data from January are available at the time of this study.




Figure 4: Boxplots per month of the hourly average wait time for all passengers across all terminals from 2013 to 2019




Figure 5: Boxplots per month of the number of hourly arriving passenger across all terminals from 2013 to 2019




Figure 6: Yearly comparison of the average wait time for all passengers per hour across all airports

Figure 7 shows how long passengers typically wait per day of the week. No clear trend is distinguishable for a particular day of the week.




Figure 7: Average wait time distribution per day of the week from 2013 to 2019

C. Wait time differences

The trends presented so far were for any passenger. Yet, U.S. citizens and foreigners enter the U.S. through separate lines, and statistics on wait times are available for each category specifically. For the same number of passengers, processing foreigners at a booth typically takes longer, since it requires checking more paperwork, such as supporting entrance documents and visas, whereas U.S. citizens only need to present their passports. The average wait time for all passengers over the year 2017 is 16.7 minutes and 16.2 minutes over the year 2018. Figure 8 shows the break down by U.S. passengers and non-U.S. passengers for these two years. Non-U.S. passengers on average spend twice as much time in line at immigration, and experience higher volatility in wait times.

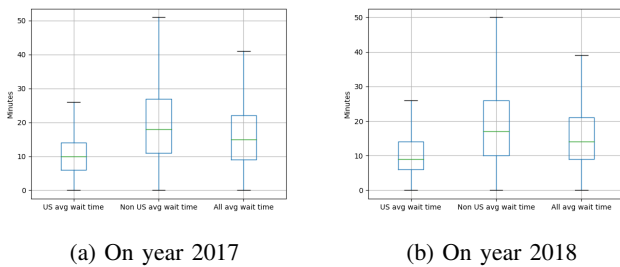


Figure 8: Average wait time distribution for passengers from January to December for the years 2017 and 2018.

Figure 9 shows the different wait times distribution per year for US citizens and non-US citizens from 2013 to 2019. The average wait time for non-US citizens does not vary much over these years, while being twice as more important as for US citizens throughout these years. Though 2019 seems to be better, observations from Section II-B indicates that the month of January is not representative of the yearly distribution.

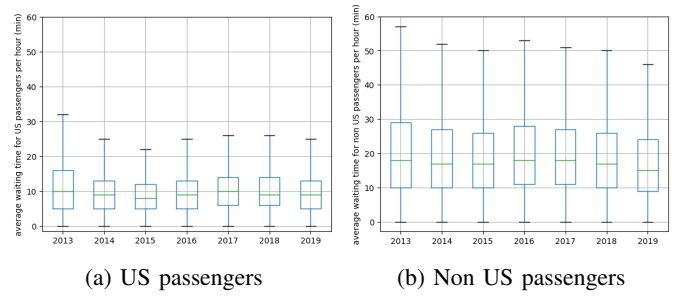


Figure 9: Average wait time and standard deviation for passengers from 2013 to 2019 across all airports

III. AIRPORTS COMPARISON

A. Overall comparison

Figure 10a shows the yearly box-plots per airports of the average wait time for all passengers for the year 2018. Interestingly both Orlando airports in Florida (SFB and MCO) show particularly high wait times with high volatility. Second in line for overall longest wait times is Hawai (SPN), before airports from the San Francisco area (SMF and SFO). Ontario International airport (ONT) is not considered in this top five due to its low volatility compared to SFO.

This high wait time is not necessarily due to lack of means. Figure 10b shows the distribution of the number of open booths per airport for the year 2018. SFO has one of the highest median of open booths, which contrasts with its wait time performance.

Figure 10c shows the distribution of the number of arriving passengers per airport for the year 2018. A first observation is that airports with high passenger volume volatility have in general a high flexibility regarding the number of open booths.


A comparison of the ratio of arriving passengers and open booths will appear in a follow-up publication.

B. Worst and Best case scenarii


From Figure 4 & 5, one can infer that the worst month for entering the United States is August: the highest median average wait time with high volatility combined with one of the largest distribution of arriving passengers. Figure 11a shows the distribution of the average wait times for all passengers per airports in August over the last six years. Airports ONT and CVG were removed since there is only two years of data for these two airports. The same two airports from Orlando Florida appear in the top five worst performing airports. On the West coast, the previously spotted airports are joined by OAK and FAT join Orlando in the top six worst performing airports.

The performance of the first four airports are even worse when considering the number of arriving passengers shown in Figure 11b. They do not have the highest median number of arriving passenger per hour nor the highest volatility.


On the other hand, from Figure 4 & 5, February appears like the best month to enter the US: lowest average wait time and volatility as well as the lowest volume of arriving passengers.



(a) Comparison of the average wait time per hour across airports over the year 2018



(b) Comparison of the number of open booths per hour across airports over the year 2018




(c) Comparison of the number of arriving passengers per hour across airports over the year 2018


Figure 10: Airport comparison using boxplots over the year 2018

Figure 12a shows the distribution of the average wait times for all passengers per airports in February over the last six years. Palm Beach International airport (PBI) and Chicago Midwest International airport (MDW) have the lowest hourly wait times on average but with a comparatively high volatility. Second best in line are Charlotte Douglas International airport (CLT) and Baltimore/Washington International airport (BWI) with a combination of low average and low volatility.

When combining these observations with the number of arriving passengers distribution shown in Figure 12b, CLT has more merit seeing how few passengers enter through PBI during that month of the year.



(a) Comparison of the average wait time per hour across airports



(b) Comparison of the number of arriving passengers per hour across airports


Figure 11: Airport comparison using boxplots for the month of August from 2013 to 2019

C. US vs. non-US wait times


Figure 13a & 13b show the distribution of the average wait times for US and non-US passengers per airports during the year 2018. The median average wait time for US citizens is reported in Figure 13b for a better comparison between the two categories. A first observation is that airports with short wait times for US citizens do not necessarily shorter wait times for non-US citizens than airports with high wait times for US citizens. For example, SPN has the second worst wait time for non US passengers while having the third shortest wait time for US passengers. Only two airports (STL and PBI) have a lower median average wait time for non-US passengers than for US passengers. Figure 13c, which shows the distribution of the average wait time ratio of non-US citizens over US citizens, indicates that only five airports have a median ratio outside the range [1,2.5]: STL and PBI with a ratio lower than one as noted previously, along with MCO, Guam airport (GUM) and SPN with ratios greater than 2.5.

IV. HOURLY WAIT TIME PREDICTION ACROSS AIRPORTS

Typical modeling of queues at airports relies on queuing theory studied in Operations Research to evaluate queue length and service time [14]. In this paper, we choose to adopt a different approach. Leveraging machine learning techniques, our goal is to predict the average wait time per hour at any



(a) Comparison of the average wait time per hour across airports



(b) Comparison of the number of arriving passengers per hour across airports

Figure 12: Airport comparison using boxplots for the month of February from 2013 to 2019

airport. ONT and CVG having less data than the other airports, they were not considered in the following study.


A. Features and regressors

The problem of interest falls in the category of regression techniques [15], more specifically under time series forecasting. The data set is partitioned into a train set and a test set. Each row in the data set corresponds to a particular hour, and the corresponding label is the average wait time for all passengers for the next hour. The following base set of features is created:


- Month of the year,
- Day of the month,
- Day of the week,
- Hour of the day,
- Number of passengers at this hour,
- Number of open booths at this hour,
- Number of flights at this hour.
- Ratio of passengers per open booth

Intuitively the waiting time at border security seems to depend on the state of the border area in the previous hour as well, e.g. if not all previously arrived passengers were processed, therefore the following features can be added to the base set:


- Number of passengers at the previous hour,
- Number of open booths at the previous hour,
- Number of flights at the previous hour,



(a) Comparison of the average wait time for US citizens per hour across airports



(b) Comparison of the average wait time for non-US citizens per hour across airports. In red is indicated the median wait time for US citizens from Figure 13a.



(c) Comparison of the ratio of average wait time for non-US citizens over US citizens. In red is indicated the interval [1, 2.5] for a better visualization.

Figure 13: US vs non-US wait times comparison using boxplots over the year 2018

- Ratio of passengers per open booth at the previous hour.

To avoid data leakage from the train set to the test set, we do not randomly assign data to either the train or test set, but select a time period for the train set and only assign a later time period to the test set. The models are trained on data from 2013 to 2017 and tested on data from the year 2018.

The different regression models are implemented using Python as the programming language and using the Scikit-learn library [16].

To obtain the best possible performance, we experiment with various algorithms, each having its specific advantages

and drawbacks. Below is a brief overview of each algorithm tested:

- 1) **Linear Regression** assumes that the relationship between the input variables and the measured variable is linear with some noise, and estimates the parameter vector with ordinary least squares minimization. It is the simplest regression method and is easily interpretable.
- 2) **Ridge Regression**, also known as Tikhonov regularization, extends the ordinary linear regression with a penalty term in the objective function proportional to the error norm. This improves the conditioning of the problem and reduces overfitting.
- 3) **Lasso Regression** performs a linear regression with regularization as well as a variable selection.
- 4) **Random Forest Regression** is an ensemble technique that fits a number of classifying decision trees on various sub-samples of the dataset and uses averaging to improve the predictive accuracy and control overfitting.
- 5) **Gradient Boosted Tree Regression** is an ensemble technique relying on decision trees as weak learners and resampling the training samples while assigning weights to the samples. It optimizes a cost function over function space by iteratively choosing a function that points in the negative gradient direction. It is typically more robust than basis learners.
- 6) **AdaBoost** is also an ensemble of decision trees relying on boosting. However, AdaBoost and Gradient Boosting optimize different loss functions.

B. Performance measures

In order to measure the quality of the predictions, two different performance measures were computed: the R^2 score and the mean absolute error (MAE).

The R^2 score, also known as the coefficient of determination, is defined as the unity minus the ratio of the residual sum of squares over the total sum of squares: $R^2 = 1 - \frac{\sum_i (y_i - f_i)^2}{\sum_i (y_i - \bar{y})^2}$, where y is the value to be predicted, \bar{y} its mean and f is the predicted value. It ranges from $-\infty$ to 1, 1 being a perfect prediction and 0 meaning that the prediction does as well as constantly predicting the mean value for each occurrence. In the case of a negative R^2 , then the model has a worse prediction than if it were predicting the mean value for each occurrence and therefore yields no useful predictions. The closer the R^2 score is to 1, the more accurate the prediction is.

Regarding the mean absolute error, the smaller its value is, the more accurate the prediction is. It is calculated using the following formula: $MAE = \frac{1}{n} \sum_i |f_i - y_i|$ where n is the number of values being predicted.

C. Simple prediction benchmarks

Four different simple prediction benchmarks were tested in order to have a better understanding of the ease of prediction. For the first three, the prediction consisted in taking the value from the previous year, the previous day or the previous hour. The last one consists in constantly predicting the mean value of the the training set.

As a more elaborate comparison benchmark, Facebook's time-series forecasting tool Prophet [17] was trained on the actual training set and its performance was measured on the data from the year 2018. The Prophet tool is based on an additive model where non-linear trends are fit with yearly, weekly, and daily seasonality [18]. It is described as robust to outliers and missing data with no parameter tuning necessary, therefore the default parameters of the Prophet tool were used for this forecasting benchmark.

D. Performance analysis

1) *One-hot encoding analysis:* A single model for all terminals was created by adding one-hot encoding features for the different terminals, i.e. 59 binary features were added, each one indicating whether a specific terminal is considered or not. This single model was trained on the data from 2013 to 2017 and tested on the data from 2018. 59 other models, one for each terminal, were trained and tested on the same data filtered by terminal. Figure 14 plots the MAE per terminal for each considered regressor. It shows that for each regressor the single terminal method outperforms the one-hot encoding method for a majority of terminals. For example, in the case of the Gradient Boosting regressor, having a different regressor per arriving terminal is better than the one-hot encoding method for forty terminals out of fifty-nine.






Figure 14: Mean absolute error comparison between one-hot encoding for airports and different regressors per airports

2) *Benchmark comparison:* The comparison with the chosen benchmarks was done using the single-terminal models, i.e. for each regressor type, one model was trained per arrival terminal. The R^2 score and mean absolute error of these models are aggregated in boxplots presented in Figure 15 alongside the boxplot performance of the five benchmarks.

Figure 15a shows that the Gradient Boosting regressors have the best R^2 performance, its R^2 scores being greater than 0 for 43 terminals out of 59, which is to be compared to 28 out of 59 for the first benchmark, the Prophet tool. Gradient Boosting has also the smallest R^2 performance deviation of all tested models. Though Ada Boost has the highest R^2 score of all models for one terminal (SFB Terminal A), of the six



(a) R^2 score




(b) Mean absolute error

Figure 15: Box plot comparison of the performance measures for the five chosen benchmarks and for the six chosen regressors when predicting the average wait time for all passengers


chosen regressors, it is the only model with a median R^2 score less than 0, along with the five benchmarks. The three linear models (Linear, Lasso and Ridge) have similar performances in this study, implying that the overfitting methods added in Lasso and Ridge are not necessarily relevant here. The maximum R^2 difference between Lasso and Linear or Ridge is about 0.64 while this distance is about 0.005 between Ridge and Linear.

Figure 15b, showing the MAE distribution, has a different ranking with respect to the linear regressors. With this performance measure, Lasso yields better results than Linear and Ridge. Otherwise the conclusions for this performance measure are the same as for the R^2 score. This performance measure has however a more tangible interpretation: Of all the models tested, only the Gradient Boosting regressors predict the average waiting time with an average error of 10 minutes or less for every terminal of entry. The Gradient Boosting regressors have a MAE of 5 minutes or less for 28 terminals.

3) *Training size analysis:* In order to visualize the impact of the training set's size, the regressors were trained using training sets of different lengths (i.e. different starting dates) and their performance was tested on the data from the year 2018. As shown in Figure 16, decreasing the size of the training set does not have a major impact on the performance of the Random Forest regressor and the Gradient



(a) Median mean squared error



(b) Average mean squared error

Figure 16: Evolution of the median and average regressors performance with the beginning of the training set

Boosting regressor. For the remaining regressors, the average and median performances improve when the training set size decreases. Their performance is however still not better than nor comparable with the Gradient Boosting regressor trained on the full training set.




Figure 17: Comparison per terminal of entry of the regressors' mean absolute error with the average wait time standard deviation over the year 2018

4) *Comparison with standard deviation*: Figure 17 shows a comparison per terminal of entry between four of the best performing models' mean absolute error and the standard deviation of the average wait time over the test year 2018. This plot shows that except for one exception (SNA Terminal C), the Gradient Boosting models mean absolute errors are significantly better than the standard deviation of the values to predict.

V. CONCLUSION

This paper is a first step in a systematic analysis of passenger wait times at customs across all airports of entry in the United States using publicly available data. This analysis makes it easier to uncover some long-term trends, i.e. an increase in the number of arriving flights coupled with a decrease in the number of open custom booths, while also enabling a per airport comparison. This analysis also laid the ground to implementing machine learning regression models in order to predict the average wait time per airport of entry. These models could be used to better anticipate the number of required booths once the number of arriving passengers is known.

ACKNOWLEDGMENTS

The authors would like to thank Nikunj Oza from NASA-Ames, the BDAI team from Verizon, Palo Alto as well as the French government for their financial support.

REFERENCES

- [1] Bureau of Transportation Statistics, "Bureau of Transportation Statistics, About BTS." [Online]. Available: <http://www.rita.dot.gov/bts/about>
- [2] NextGen Integration and Implementation Office, "NextGen Implementation Plan," in *Federal Aviation Administration*, 2009.
- [3] P. Ky and B. Miaillier, "SESAR: towards the new generation of air traffic management systems in europe," *Journal of Air Traffic Control*, vol. 48, no. 1, pp. 11–14, 2006.
- [4] E. Torres, J. Domínguez, L. Valdés, and R. Aza, "Passenger waiting time in an airport and expenditure carried out in the commercial area," *Journal of Air Transport Management*, vol. 11, no. 6, pp. 363–367, Nov. 2005.
- [5] A. Pruyn and A. Smids, "Effects of waiting on the satisfaction with the service: Beyond objective time measures," *International Journal of Research in Marketing*, vol. 15, no. 4, pp. 321–334, Oct. 1998.
- [6] H. Nikoue, A. Marzuoli, J.-P. Clarke, E. Feron, and J. Peters, "Passenger Flow Predictions at Sydney International Airport: A Data-Driven Queuing Approach," *arXiv:1508.04839 [cs]*, Aug. 2015.
- [7] K. Gkritza, D. Niemeier, and F. Mannering, "Airport security screening and changing passenger satisfaction: An exploratory assessment," *Journal of Air Transport Management*, vol. 12, no. 5, pp. 213–219, Sep. 2006.
- [8] D. Wang, "Methods for analysis of passenger trip performance in a complex networked transportation system," Ph.D. dissertation, George Mason University, 2007.
- [9] B. Roberts, S. McGonegal, F. Prager, D. Wei, A. Rose, C. Baschnagel, T. Beggs, and O. Baghelai, "Analysis of primary inspection wait time at U.S. ports of entry," 2014.
- [10] H. B. Sankaranarayanan, G. Agarwal, and V. Rathod, "An exploratory data analysis of airport wait times using big data visualisation techniques," in *2016 International Conference on Computation System and Information Technology for Sustainable Solutions (CSITSS)*. Bengaluru, India: IEEE, Oct. 2016, pp. 324–329.
- [11] M. Johnstone, "A dynamic architecture for increased passenger queue model fidelity," in *Proceedings of the 2009 Winter Simulation Conference (WSC)*. Austin, TX, USA: IEEE, Dec. 2009, pp. 3129–3139.
- [12] C. Robertson, S. Shrader, D. Pendergraft, L. Johnson, and K. Silbert, "The role of modeling demand in process re-engineering," in *Proceedings of the Winter Simulation Conference*, vol. 2. San Diego, CA, USA: IEEE, 2002, pp. 1454–1458.
- [13] U. S. Customs and B. Protection, "Airport wait times." [Online]. Available: <https://awt.cbp.gov>
- [14] L. Kleinrock, *Queueing systems, volume 2: Computer applications*. Wiley New York, 1976, vol. 66.
- [15] C. M. Bishop, *Pattern recognition and machine learning*. Springer, 2006.
- [16] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, and D. Cournapeau, "Scikit-learn: Machine Learning in Python," *Machine Learning in Python*, 2011.
- [17] S. J. Taylor and B. Letham, "Forecasting at Scale," *The American Statistician*, vol. 72, no. 1, pp. 37–45, Jan. 2018.
- [18] Facebook, "Prophet - Forecasting at scale." [Online]. Available: <https://facebook.github.io/prophet/>

APPENDIX

Table I: Terminal of arrivals abbreviations

Abbreviation	Airport (IATA)	Terminal name
ATL - CE	ATL	Concourse E
ATL - CF	ATL	Concourse F
AUS - M	AUS	Main
BOS - TE	BOS	Terminal E
BWI - IA	BWI	International Arrivals
CLT - M	CLT	Main
DEN - I	DEN	International
DFW - TD	DFW	Terminal D
DTW - MT	DTW	McNamara Terminal
DTW - NT	DTW	North Terminal
EWB - TB	EWB	Terminal B
EWB - TC	EWB	Terminal C
FAT - M	FAT	Main
FLL - T1	FLL	Terminal 1
FLL - T4	FLL	Terminal 4
GUM - MT	GUM	Main Terminal
HNL - MOT	HNL	Main Overseas Terminal
IAD - IA	IAD	International A
IAH - I	IAH	IAB
JFK - T1	JFK	Terminal 1
JFK - T4	JFK	Terminal 4 (IAT)
JFK - T5B	JFK	Terminal 5 (Jet Blue)
JFK - T7	JFK	Terminal 7 (British)
JFK - T8	JFK	Terminal 8 (American)
LAS - T3	LAS	Terminal 3
LAX - S2	LAX	Satellite 2
LAX - S5	LAX	Satellite 5
LAX - S7	LAX	Satellite 7
LAX - T4	LAX	Terminal 4
LAX - TBIT	LAX	Tom Bradley International Terminal
MCO - A1	MCO	Airside 1
MCO - A4	MCO	Airside 4
MDW - MT	MDW	Main Terminal
MIA - CT	MIA	Central Terminal
MIA - NT	MIA	North Terminal
MIA - ST	MIA	South Terminal
MSP - T1L	MSP	Terminal 1 Lindbergh
MSP - T2H	MSP	Terminal 2 Humphrey
OAK - M	OAK	Main
ORD - T5	ORD	Terminal 5
PBI - M	PBI	Main
PDX - M	PDX	Main
PHL - TA	PHL	Terminal A
PHX - M	PHX	Main
RDU - T2	RDU	Terminal 2
SAN - M	SAN	Main
SAT - M	SAT	Main
SEA - SS	SEA	South Satellite
SFB - TA	SFB	Terminal A
SFO - TA	SFO	Terminal A
SFO - TG	SFO	Terminal G
SJC - M	SJC	Main
SJU - SJA	SJU	San Juan AA
SLC - M	SLC	Main
SMF - MT	SMF	Main Terminal
SNA - TC	SNA	Terminal C
SPN - IA	SPN	International Arrivals
STL - M	STL	Main
TPA - M	TPA	Main