Passengers on social media: A real-time estimator of the state of the US air transportation system

Philippe Monmousseau, Aude Marzuoli, Eric Féron, Daniel Delahaye

► To cite this version:

Philippe Monmousseau, Aude Marzuoli, Eric Féron, Daniel Delahaye. Passengers on social media: A real-time estimator of the state of the US air transportation system. EIWAC 2019:, 6th ENRI International Workshop on ATM/CNS, ENRI, Oct 2019, Tokyo, Japan. pp 189-205, 10.1007/978-981-33-4669-7_11 . hal-02364816

HAL Id: hal-02364816

<https://hal-enac.archives-ouvertes.fr/hal-02364816>

Submitted on 15 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.

[EN-A-46] Passengers on social media: A real-time estimator of the state of the US air transportation system

(EIWAC 2019)

⁺P. Monmousseau*, A. Marzuoli**, E. Feron**, D. Delahaye*

* Ecole Nationale de l'Aviation Civile (ENAC)

Universite de Toulouse

Toulouse, France

[philippe.monmousseau | delahaye]@recherche.enac.fr

**School of Aerospace Engineering

Georgia Institute of Technology

Atlanta GA, USA

[amarzuoli3 | feron]@gatech.edu

Abstract: This paper aims at investigating further the use of the social media Twitter as a real-time estimator of the US Air Transportation system. Two different machine learning regressors have been trained on this 2017 passenger-centric dataset and tested on the first two months of 2018 for the estimation of air traffic delays at departure and arrival at 34 different US airports. Using three different levels of content-related features created from the flow of social media posts led to the extraction of useful information about the current state of the air traffic system. The resulting methods yield higher estimation performances than traditional state-of-the-art and off-the-shelf time-series forecasting techniques performed on flight-centric data for more than 28 airports. Moreover the features extracted can also be used to start a passenger-centric analysis of the Air Transportation system. This paper is the continuation of previous works focusing on estimating air traffic delays leveraging a real-time publicly available passenger-centered data source. The results of this study suggest a method to use passenger-centric data-sources as an estimator of the current state of the different actors of the air transportation system in real-time.

Keywords: delay estimation, ATM performance measurement, big data, machine learning

1. INTRODUCTION

The Air Transportation System is a complex interconnected system that carried more than 631 million passengers on domestic flights in the United States in 2010 according to the Bureau of Transportation Statistics (BTS) [1]. Flight delays are still a major issue both in the United States with 27.0% of departing flights and 27.8% of arriving flights experiencing delays in 2017 [1].

Most previous studies aimed at predicting or classifying flight delays were centered on flight-centric information coming from a variety of sources with different levels of public availability, and using only very little passenger-centric data. Mueller and Chatterji [2] created a probabilistic model of delays by fitting Poisson and Normal distributions to the historic delay data from 10 airports. Rebollo and Balakrishnan [3] implemented a network model to classify and predict future delays on specific links or specific airports using two years of flight-centric and weather-related data. Klein et al. [4] and [5] focused on predicting short-term weather-related

delays using only past and current weather information. Aljubairy et al. [6] used Internet of Things in order to analyze flight-related sensors in real-time and classify the delay of an upcoming flight.

Over the past few years, NextGen [7] in the United States has been advocating a shift from flight-centric metrics to passenger-centric metrics to evaluate the performance of the Air Transportation System. The failures and inefficiencies of the air transportation system not only have a significant economic impact but they also stress the importance of putting the passenger at the core of the system [8] [9]. Several studies have highlighted the disproportionate impact of airside disruptions on passenger door-to-door journeys. Flight delays do not accurately reflect the delays imposed upon passengers' full multi-modal itinerary. Cook et al. [10] designed propagation-centric and passenger-centric performance metrics, and compare them with existing flight-centric metrics. Wang [11] showed that high passenger trip delays are disproportionately generated by canceled flights

and missed connections. Nine of the busiest thirty-five airports cause 50% of total passenger trip delays. Congestion, flight delay, load factor, flight cancellation time and airline cooperation policy are the most significant factors affecting total passenger trip delay. NextGen intends to not only improve the predictability and resilience of the US Air Transportation System, but also to reduce door-to-door travel time for passengers.

Passengers are at the core of this system and, yet, limited quantitative information about passenger movements is publicly shared. Each aviation stakeholder only has access to a partial view of the passenger-side of air transportation operations making a system-wide data-driven picture of passenger behavior difficult to implement. The BTS provides aggregated passenger data per market but no granular information. Passenger surveys conducted by airports or airlines, while very detailed, remain limited to small samples of passengers and short time periods, and may not be representative.

Precursor work was made by Marzuoli et al. in [12] and [13] using mobile phone data in order to analyze the performances of airports from the passengers' perspective. These studies validated the use of passenger-centric data to better assess the overall health of the Air Transportation System. However mobile phone data is proprietary data and is not often publicly available. In order to operate in real-time, it is thus necessary to also look into other sources of passenger data available on a national scale.

Another popular source of data previously used for studying large-scale behaviors with real time availability is social media, in particular Twitter. With more than 68 millions active users in the United States [14], Twitter is an important pool of user-created data that is still not fully leveraged. Twitter has already been the main focus of many studies focused on its real-time availability, especially during natural disasters with multiple works by Palen et al. [15], [16], [17]. Terpstra et al. also studied how a real time Twitter analysis could have provided valuable information for the operational response of a natural disaster crisis management with the case of the storm hitting a festival in Belgium [18]. Regarding the air transportation field, most works mining Twitter data focus on how airlines are perceived by passengers by means of sentiment analysis [19] or sentiment classification [20]. Though these works give a good insight on how passengers perceive the state of some specific actors within the air transportation system, it does not give a global idea of its health. Monmousseau et al. in [21] used publicly available social media data created by passengers to accurately estimate and predict the hourly aggregated status of the US air transportation system.

This paper proposes to build on this previous work in order to estimate the state of the air transportation

system to a finer level. Rather than predicting the number of delays across all the United States, the proposed passenger-centric models are improved and tuned to accurately estimate the state of delays for each of the 35 major airports within the United States. The created models are based on three different levels of content-related features created from the flow of social media posts. First results indicate that these new models can estimate the number of hourly delays with a mean absolute error of less than 3 flights for 26 of the considered airports, and of less than 6 flights for the 9 remaining airports.

The rest of the paper is structured as follows: Section 2 describes the datasets and the feature extraction process. The methodology and results of the training process are shown in Section 3, before being analyzed and exploited in Section 4. Section 5 concludes this study and discusses possible future steps.

2. DATASET DESCRIPTION AND FEATURE SELECTION

2.1 Dataset description

Following the initial work performed in [21], the goal here is to use passengers behavior on social media - in particular on Twitter - in order to analyze and estimate the flight-centric health of the US air-transportation system at an airport level. In this study, the flight-centric health of an airport is described by delay related information contained within BTS data. This data is publicly available usually with a two to three month delay and this study limits itself with the BTS data from January 2017 to February 2018.

The Twitter dataset available for this study is the same as in [21] and consists of all the tweets found using a basic search for each handle of 7 major US airlines as well as 34 major US airports (one of them having two Twitter handles). The full list of handles can be found in Table 1. Each entry consists of a timestamp, a user id, the content of the tweet and the handle used to retrieve the tweet. This dataset spans the entire period from January 1st 2017 to February 28th 2018. The extraction of features from this dataset has been improved since the previous study and is described in Section 2.2.

Figure 1 plots the total number of tweets related to each handle over the year 2017 against the total number of flights flown by each airline or to and from each airport. Airlines tend to gather more tweets than airports, and the number of tweets is not necessarily correlated to the number of flights flown per airline. The handle "@Delta" gathered the most tweets over 2017 even though Southwest Airlines carried out the most flights in 2017. Most airports are regrouped around a cluster of 10k tweets and 200k flights over 2017, with

Table 1: Twitter handles used for gathering tweets

Category	Twitter handles
Airlines	@united, @Delta, @AmericanAir, @SouthwestAir, @SpiritAirlines, @VirginAmerica, @JetBlue
Airports	@JFKairport, @ATLairport, @flyLAXairport, @fly2ohare, @DFWairport, @DENairport, @CLTairport, @LASairport, @PHXSkyHarbor, @MiamiAirportMIA, @iah, @EWRairport, @MCOairport, @Official_MCO, @SeaTacAirport, @mspairport, @DTweetin, @BostonLogan, @PHLairport, @LGAairport, @FLLFlyer, @BWI_Airport, @Dulles_Airport, @MidwayAirport, @Reagan_Airport, @slcaairport, @SanDiegoAirport, @flyTPA, @flypdx, @flystl, @flySFO, @HobbyAirport, @flynashville, @AUSTinAirport, @KCIairport

Los Angeles International airport (LAX) and Hartsfield-Jackson Atlanta International airport (ATL) being exceptions due to their higher number of tweets.

Figure 1: Number of tweets over the number of flights during the year 2017 for the considered airlines and airports

In order to estimate the flight-centric health of each considered airport, this information first needs to be extracted from the BTS dataset for each airport. Only two types of delayed flights are considered here from a passenger’s perspective: Flights departing with any amount of delay, and flights arriving with a delay greater than 15 minutes. Once all the flights departing an airport and all the flights arriving at the same airport are selected, the following values can be aggregated per hour:

- NumDepDelay: Number of flights departing with a delay
- NumArrDelay15: Number of flights arriving with a delay greater than 15 minutes

The aim of this study is to accurately estimate these two values for each airport at every hour using a single passenger-centric dataset.

2.2 Feature selection on Twitter data

2.2.1 Volume features

Features were extracted identically for all search handles presented in Table 1, for the exception of @MiamiAirportMIA, which does not gather enough tweets. In addition to the raw number of tweets per hour per search handle, keyword related information is also extracted

from the Twitter dataset. In order to keep all the relevant tweets without having to decline all the possible forms of the chosen keywords (e.g. "delay", "delayed", "delays", etc.), simple regular expression filters were created for each keyword: Any tweet containing a word starting with the related keyword is kept and the resulting tweets are then aggregated per hour. Five keywords were chosen for this study: 'delay', 'wait', 'cancel', 'hours', 'refund'.

2.2.2 Topic features

Another way of exploiting information from the content of these tweets is to perform a topic analysis of the tweet database using Latent Dirichlet Allocation [22] (LDA). A first step in topic analysis is to clean the documents analyzed, here the tweets. This cleaning process was already performed in [13] and [21] and consists of the following steps: Any reference to websites or pictures was replaced by a corresponding keyword. Every mention to another Twitter user within a tweet (@someone) as well as most emojis were similarly replaced. Since this database contains many replies from airlines to their customers, individual signatures of each agent were also replaced by a keyword. Dates and times were also generically replaced by keywords (e.g. "3rd Jan 2017" becomes "DATE" and "4pm" becomes "TIME"). The resulting text was then filtered from common stop-words and from words occurring only once in the whole year of 2017. A list of 100 topics was then created using the Gensim [23] library. The topic distribution of each tweet was then calculated before averaging these distributions per hour and per search handle. The hourly standard deviation of the distributions was also extracted.

This cleaning process introduces two additional keywords that enables a quick filtering of tweets, and therefore two additional features to add per search handle: tweets containing a picture and those containing a website link. Thus seven keywords are actually considered for feature extraction: 'delay', 'wait', 'cancel', 'hours', 'refund', 'PICTURE', 'WEBSITE'.

2.2.3 Sentiment features

Sentiment analysis is also used here to enhance the feature set considered. Two different datasets and cleaning method were used to train three different regressors each. The first dataset used was the labelled dataset used in a Kaggle competition [24] and was cleaned using the same process as for the previous LDA learning. The generic

keywords from the cleaning process (e.g. 'WEBSITE', 'DATE') were removed before creating the associated dictionary, as well as words appearing in less than 20 tweets or in more than 75% of the full dataset. A second dataset and cleaning process was generated based on the work of Read [25]. Emoji filters were used to extract tweets from the initial dataset and automatically label them with a positive or negative sentiment according to Table 2. The text cleaning process is improved by merging negation words ("no", "not" and "never") with the word that follows it. The tokens used for the creation of the dictionary are the resulting bigrams, i.e. combinations of two words that follow each other in a tweet, with the same frequency filter as the first method described.

Table 2: Emoji sentiment association

Category	Emojis
Positive	":)", "=)", ":-)", ":-)", ":-)", ":-D", ":-D", "=D"
Negative	":(", ":-(", "=(", ":-@", ":-(", ":- "

For both methods, three classifiers are trained (a random forest classifier, a naive Bayesian classifier and a logistic regressor) using the scikit-learn library [26]. A sentiment score is then calculated for each tweet by averaging the output of these classifiers, 0 meaning a unanimous negative sentiment and 1 a unanimous positive sentiment. The hourly average of these scores are added to the Twitter feature set.

2.2.4 Summary

Given the temporal nature of the data analyzed, the following features were chosen to keep track of the date: month of the year, day of the month, day of the week and hour in the day. In summary the following 8,484 features are considered:

- Hourly volume of tweets for each search handle (7 airlines and 33 airports giving 40 features): `Num_tweets_handle`
- Hourly volume of keyword-related tweets for each search handle (40x7 features): `Num_tweets_keyword_handle`
- Hourly average of tweets' sentiment (40x2 features): `Mean_sent_method_handle`
- Hourly average of topic distribution for each search handle (40x100 features): `Mean_topic_handle`
- Hourly standard deviation of topic distribution for each search handle (40x100 features): `Std_topic_handle`
- Month of the year, Day of the month, Day of the week and Hour in the day (4 features)

3. ESTIMATING DELAYS

The aim of this section is to see how well it is possible to estimate per airport the number of flights departing with a delay and the number of flights arriving with

a delay greater than 15 minutes using the features extracted from the Twitter dataset. The dataset was split into a training set consisting of the data from the year 2017, and a testing set with the data from January and February 2018.

3.1 Methodology

For each BTS value, two different machine learning regressors were trained on the training data set: a Random Forest regressor and a Gradient Boosting regressor. These regressors were implemented from scikit-learn [26] with identical hyper-parameters. The maximum depth of each regressor was limited to ten, the minimum number of samples for a split was fixed to two and the maximum number of trees was fixed at ten.

As a comparison benchmark, we used Facebook's time-series forecasting tool Prophet [27] on the 2017 BTS data to forecast the full two first months of 2018. The Prophet tool is based on an additive model where non-linear trends are fit with yearly, weekly, and daily seasonality [28]. It is described as robust to outliers and missing data with no parameter tuning necessary, therefore the default parameters of the Prophet tool was used for this forecasting benchmark.

Lastly, the standard deviation of the BTS values in the training set were calculated to illustrate the added value of the trained regressors. The performance measures used to compare the different regressors are presented in the upcoming section 3.2.

3.2 Estimation performance measures

In order to measure the performance of the different models, two different indicators were used: 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:

$$R^2 = 1 - \frac{\sum_i (y_i - f_i)^2}{\sum_i (y_i - \bar{y})^2} \quad (1)$$

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.

Regarding the mean-absolute error, the smaller its value is, the more accurate the prediction is. It is calculated using the following formula:

$$\text{MAE} = \frac{1}{n} \sum_i |f_i - y_i| \quad (2)$$

where n is the number of values being predicted.

3.3 Estimation results

Fig.2 shows a comparison per airport of the mean-absolute error of the two trained regressors along with the chosen benchmark for the estimation of the number of flights departing with a delay. The standard deviation of the number of delayed departing flights at each airport during the year 2017 is also included for comparison. The Random Forest models have the best results in this case: they outperform the Gradient Boosting models at all-but-one airports (LAX) and the Facebook Prophet tool on 31 airports out of 34. For 26 airports, the Random Forest models are able to estimate the hourly number of delayed departing flights with a mean-absolute error of three flights or less, and with an error of less than six flights for the remaining airports.

Figure 2: Comparison of the mean absolute errors per airport for the trained regressors for the estimation of the number of delayed departing flights. The standard deviation of the BTS value on the training set is included for comparison.

Fig.3 shows a comparison per airport of the mean-absolute error of the two trained regressors along with the chosen benchmark for the estimation of the number of flights arriving with a delay greater than 15 minutes. The standard deviation of the number of delayed arriving flights at each airport during the year 2017 is also included for comparison. The Random Forest models also have the best results in this case though their relative performance are not as important as for delayed departing flights : they outperform the Gradient Boosting models at 27 airports out of 34 and the Facebook Prophet tool on 28 airports out of 34. The absolute performance is however better than for estimating the number of delayed departing flights. For 28 airports, the Random Forest models are able to estimate the hourly number of delayed departing flights with a mean-absolute error of less than three flights, and with an error of less than five flights for the remaining airports.

Fig. 4 shows a comparison per airport of the R^2 score of the two trained regressors along with the chosen benchmark for the estimation of the number of flights

Figure 3: Comparison of the mean absolute errors per airport for the trained regressors for the estimation of the number of flights arriving with a delay greater than 15 minutes. The standard deviation of the BTS value on the training set is included for comparison.

departing with a delay. The Random Forest models still have the best results in this case, but the model associated with LAX airport also shows the only negative score. They outperform the Gradient Boosting models at 27 airports out of 34 and the Facebook Prophet tool on 28 airports out of 34.

Figure 4: Comparison of the R^2 scores per airport for the trained regressors for the estimation of the number of delayed departing flights

4. ANALYSIS AND APPLICATIONS

The aim of this section is to analyze the differences between the chosen models as well as explore possible applications resulting from the extracted features.

4.1 Model analysis

Figure5 shows the hourly prediction of the number of delayed departing flights at Atlanta airport (ATL) over the period January 12th-16th for the two trained regressors along with the benchmark and the actual values. This airport was chosen since it has the highest BTS standard deviation for the number of delayed

departing and arriving flights, and the period was chosen to illustrate the high variability of the number of delays from a day to another. In this example, January 12 has more than twice as many delayed flights than any other day, as well as important hourly variations.

Figure 5: Predicted number of delayed departing flights at ATL by the trained regressor over the period January 12th, 2018 to January 16th, 2018. The actual number of delayed departing flights is indicated for comparison.

Figure 5 illustrates the main differences between the different models: The Prophet tool predicts for each day a similar daily variation with three peaks during the day yet with amplitudes varying depending on the month and the day of the week. It also predicts negative values, which underlines some limitations of the model in this case. The added value from passenger-centric data-sources is better seen on January 12 and 13, where only the Random Forest regressor is able to estimate the higher number of delays on January 12 before correctly estimating the more usual levels of January 13. The Gradient Boosting regressor doesn't estimate outliers as well as the Random Forest regressor due to the difference in their loss functions. That difference is also illustrated by the non-zero minimum of the Gradient Boosting estimation during night time.

4.2 Other applications

4.2.1 Real-time sentiment analysis

The extracted features can be fed to the trained models for accurately estimating the number of delayed flights, but they can also be used directly in order to sense the overall passenger mood. Once the sentiment analyses are conducted on the tweets, it is possible to merge them into one score per airline and monitor their evolution.

Figure 6 shows the hourly average mood for three major airlines during the Northeastern bomb cyclone

Figure 6: Average passenger sentiment with respect to three major airlines over the period January 2nd, 2018 to January 6th, 2018, corresponding to a bomb cyclone hitting in the North-East of the US.

studied in [13]. These three airlines have a similar passenger mood evolution at the beginning and the end of the period, yet United Airlines shows a drop in passenger mood on January 4th, the day when the bomb cyclone actually hit the East coast. Though all three airlines have hubs in New York, United Airlines is the only airline with a hub at Newark International Airport (EWR) and not John F. Kennedy International Airport (JFK) nor LaGuardia Airport (LGA), which were both closed during the bomb cyclone, meaning that United Airlines probably had more dissatisfied passengers to handle on site during these extreme weather conditions.

4.2.2 Airports passenger map

After training the Random Forest models, it is possible to search for the most important features within the 8,484 initial features for each airport. This is achieved by using the Mean Decrease Impurity measure defined by Breiman in [29] and normalizing the obtained feature importances so that the sum of all feature importances is equal to one. Table 3 shows the ten features with the highest feature importances for predicting the number of delayed departing flights in ATL. Besides date related features, four of the top ten features are related to the volume of tweets containing delay keywords.

Table 3: Top ten features for predicting the number of delayed departing flights at ATL

Rank	Feature	Rank	Feature
1	Hour	6	DayOfMonth
2	Month	7	delay_@SouthwestAir
3	DayOfWeek	8	num_ATL
4	delay_@Delta	9	delay_JFK
5	delay_ATL	10	mean_63_BWI

Once the features gathering 99% of the total importance for estimating the number of delayed flights are extracted, it is possible to group these features per origin in order to gain some insight on how airports are related from a passenger perspective. For example, once the most important features for estimating the number of delays at ATL are extracted, it is possible to count how many of these features are issued from tweets gathered using the handle of John F. Kennedy International Airport (JFK).

Figure 7: Map of feature links between Atlanta airport (ATL) and the other airports for estimating the number of delayed departing flights. The larger the link, the more features were kept among the features gathering 99% of the total importance for estimating the number of departing delayed flights at ATL.

Figure 8: Map of delay links between Atlanta airport (ATL) and the other airports. The larger the link, the more flights departed with a delay during 2017 from ATL towards the connecting airport. Only links with more than 1000 delayed flights in 2017 were considered.

Figure 7 shows how ATL is connected to the other airports from this perspective. The larger the link between ATL and another airport, the more features were kept among the features gathering 99% of the total importance for estimating the number of departing delayed flights at ATL. Interestingly, this airport graph is different from the graph built from the actual BTS values. Figure 8 shows how ATL is connected to the other airports using the number of delayed departing flights from ATL. For example, although there are many delayed flights departing to Florida, few features from Floridan airports are kept. The opposite observation can be made regarding Portland (PDX): there were less than a thousand delayed flights from ATL to PDX, yet features from PDX were kept.

This example illustrates the possibility of creating a yearly review of airport relationship from a passenger point of view. Future studies should investigate more thoroughly the possible correlation and relation between the passenger connection map and the delay connection map.

5. CONCLUSION

This paper aimed at investigating further the use of the social media Twitter as an estimator of the US Air Transportation system. Exploiting both raw volume information as well as different levels of content information within the Twitter stream enables to accurately estimate for each airport the number of flights departing with a delay and the number of flights arriving with a delay greater than fifteen minutes. This passenger-based estimation yields a better estimation performance for a majority of airports compared to using a state-of-the-art and off-the-shelf forecasting tool on the flight-centric data alone. Moreover, the methods used to extract relevant features from this passenger-centric data-source can be used to gain additional real-time insight on how passengers relate to the Air Transportation system.

This study confirmed that information contained in passenger-centric datasets are useful for a better understanding of the different stakeholders within the air transportation system, and have the added benefit of being more readily and publicly available than flight-centric datasets. Future studies should focus on analyzing cases when the estimation is less accurate, implying differences between the handling of passengers and that of planes. Another direction of study considered is to validate this method to other countries or regions (e.g. the European Union) where sufficient flight-centric data is available.

ACKNOWLEDGMENTS

The authors would like to thank Nikunj Oza from NASA-Ames, the BDAI team from Verizon Media in Sunnyvale 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] E. Mueller and G. Chatterji, "Analysis of Aircraft Arrival and Departure Delay Characteristics," in *AIAA's Aircraft Technology, Integration, and Operations (ATIO) 2002 Technical Forum*. Los Angeles, California: American Institute of Aeronautics and Astronautics, Oct. 2002.
- [3] J. J. Rebollo and H. Balakrishnan, "Characterization and prediction of air traffic delays," *Transportation Research Part C: Emerging Technologies*, vol. 44, pp. 231–241, Jul. 2014.
- [4] A. Klein, "Airport delay prediction using weather-impacted traffic index (WITI) model," in *29th Digital Avionics Systems Conference*. Salt Lake City, UT, USA: IEEE, Oct. 2010, pp. 2.B.1–1–2.B.1–13.

- [5] B. Sridhar and N. Chen, "Short-term national airspace system delay prediction using weather impacted traffic index," *Journal of guidance, control, and dynamics*, vol. 32, no. 2, pp. 657–662, 2009.
- [6] A. Aljubairy, A. Shemshadi, and Q. Z. Sheng, "Real-time investigation of flight delays based on the Internet of Things data," in *Advanced Data Mining and Applications*, J. Li, X. Li, S. Wang, J. Li, and Q. Z. Sheng, Eds. Cham: Springer International Publishing, 2016, vol. 10086, pp. 788–800.
- [7] NextGen Integration and Implementation Office, "NextGen Implementation Plan," in *Federal Aviation Administration*, 2009.
- [8] World Economic Forum, "Connected World : Transforming Travel, Transportation and Supply Chains," <http://www3.weforum.org/docs>, 2013.
- [9] —, "Smart travel: Unlocking economic growth and development through travel facilitation," <http://www3.weforum.org/docs/GAC/2014>, 2014.
- [10] A. Cook, G. Tanner, S. Cristóbal, and M. Zanin, "Passenger-Oriented Enhanced Metrics," p. 8, 2012.
- [11] D. Wang, "Methods for analysis of passenger trip performance in a complex networked transportation system," Ph.D. dissertation, George Mason University, 2007.
- [12] A. Marzuoli, E. Boidot, E. Feron, and A. Srivastava, "Implementing and validating air passenger-centric metrics using mobile phone data," *Journal of Aerospace Information Systems*, 2018.
- [13] A. Marzuoli, P. Monmousseau, and E. Feron, "Passenger-centric metrics for Air Transportation leveraging mobile phone and Twitter data," in *Data-Driven Intelligent Transportation Workshop - IEEE International Conference on Data Mining 2018*, Singapore, Nov. 2018.
- [14] Statista, "Monthly active Twitter users in the united states," <https://www.statista.com/statistics/274564/monthly-active-twitter-users-in-the-united-states/>.
- [15] L. Palen, K. Starbird, S. Vieweg, and A. Hughes, "Twitter-based information distribution during the 2009 Red River Valley flood threat," *Bulletin of the American Society for Information Science and Technology*, vol. 36, no. 5, pp. 13–17, Jun. 2010.
- [16] S. Vieweg, A. L. Hughes, K. Starbird, and L. Palen, "Microblogging during two natural hazards events: What twitter may contribute to situational awareness," p. 10, 2010.
- [17] K. Kireyev, L. Palen, and K. M. Anderson, "Applications of Topics Models to Analysis of Disaster-Related Twitter Data," p. 4, 2009.
- [18] T. Terpstra and R. Stronkman, "Towards a realtime Twitter analysis during crises for operational crisis management," p. 10, 2012.
- [19] J. O. Breen, "Mining twitter for airline consumer sentiment," *Practical text mining and statistical analysis for non-structured text data applications*, vol. 133, 2012.
- [20] Y. Wan and Q. Gao, "An ensemble sentiment classification system of Twitter data for airline services analysis," in *2015 IEEE International Conference on Data Mining Workshop (ICDMW)*. Atlantic City, NJ, USA: IEEE, Nov. 2015, pp. 1318–1325.
- [21] P. Monmousseau, A. Marzuoli, E. Feron, and D. Delahaye, "Predicting and Analyzing US Air Traffic Delays using Passenger-centric Data-sources," in *Thirteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2019)*, Vienna, Austria, 2019.
- [22] D. M. Blei, "Latent Dirichlet Allocation," *Journal of machine Learning research*, pp. 993–1022, 2003.
- [23] R. R. P. Sojka, "Software Framework for Topic Modelling with Large Corpora," p. 5.
- [24] Kaggle, "Twitter US airline sentiment," <https://www.kaggle.com/crowdflower/twitter-airline-sentiment>.
- [25] J. Read, "Using emoticons to reduce dependency in machine learning techniques for sentiment classification," in *Proceedings of the ACL Student Research Workshop on - ACL '05*. Ann Arbor, Michigan: Association for Computational Linguistics, 2005, p. 43.
- [26] 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.
- [27] S. J. Taylor and B. Letham, "Forecasting at Scale," *The American Statistician*, vol. 72, no. 1, pp. 37–45, Jan. 2018.
- [28] Facebook, "Prophet - Forecasting at scale." [Online]. Available: <https://facebook.github.io/prophet/>
- [29] L. Breiman, *Classification and regression trees*. Routledge, 2017.

COPYRIGHT

Copyright Statement

The authors confirm that they, and/or their company or institution, hold copyright of all original material included in their paper. They also confirm they have obtained permission, from the copyright holder of any third party material included in their paper, to publish it as part of their paper. The authors grant full permission for the publication and distribution of their paper as part of the EIWAC2019 proceedings or as individual off-prints from the proceedings.