APPLICATION OF DATA FUSION AND MACHINE LEARNING TO THE ANALYSIS OF THE RELEVANCY OF RECOMMENDED FLIGHT REROUTES

A Dissertation
Presented to
The Academic Faculty

by

Ghislain Dard

In Partial Fulfillment
of the Requirements for the Degree
Master of Science in the
School of Aerospace Engineering

Georgia Institute of Technology
May 2019

APPLICATION OF DATA FUSION AND MACHINE LEARNING TO THE ANALYSIS OF THE RELEVANCY OF RECOMMENDED FLIGHT REROUTES

Approved by:

Dr. Dimitri Mavris Advisor
School of Aerospace Engineering
Georgia Institute of Technology

Dr. Olivia Pinon Fischer
School of Aerospace Engineering
Georgia Institute of Technology

Mr. Mike Paglione
Tech Center
Federal Aviation Administration

Date Approved: April 16, 2019

“Never doubt that a small group of thoughtful, committed citizens can change the world; indeed, it's the only thing that ever has.”

Margaret Mead, 1901 – 1978, American anthropologist

ACKNOWLEDGEMENTS

First, I would like to thank my advisor Dr Dimitri Mavris, who gave me the opportunity to work in the Aerospace System Design Laboratory during three semesters and therefore work on projects with aeronautical companies such as this thesis with the Federal Aviation Administration. Thank you very much Doc, my experience at Georgia Tech would have been very different and probably not as interesting without you.

Then, I want to give a special and warm thank you to my committee member and technical advisor on this thesis, Olivia Pinon Fischer. Thank you very much for your commitment and advices. You have been an example of hard work and patience for me. It is easy and rewarding to work with you. Merci !

I would like to thank Mike Paglione, from the FAA, my third committee member. Thank you for guiding my work, providing data and trust me. From the FAA, I also want to thank Nelson Brown, David Chong, Anya Berges and Jonathan Rein. You have been very helpful and important to guide my work so it can provide benefits to the FAA.

To all my friends and lab mates from Georgia Tech, thank you for your friendship and for motivating me: Lydia Mangiwa, Federico Bonalumi, Gabriel Achour, Manon Huguenin, Marc-Henri Bleu Lainé, Asmae El Hirache.

Finally, I want to say another special thank you to Eugene Mangortey who has always supported me, gave me advices and who convinced me to work on this thesis. Merci mon ami. Good luck for your PhD.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF SYMBOLS AND ABBREVIATIONS

AAR	Airport Acceptance Rate
AFP	Airspace Flow Program
ARTCC	Air Route Traffic Control Center
ASDI	Aircraft Situation Display
ATCSCC	Air Traffic Control System Command Center
CASSIE	Computing Analytics and Shared Services Integrated Environment
CSV	Comma Separated Values
CTOP	Collaborative Trajectory Options Program
EDDS	En-Route Data Distribution System
ERAM	En-Route Automation Modernization
FIXM	Flight Information Exchange Model
FAA	Federal Aviation Administration
FYI	For Your Information
GADV	General Advisory
GDP	Ground Delay Programs
GS	Ground Stops
HADDS	Host Air Traffic Management (ATM) Data Distribution System
MAE	Mean Absolute Error
MINIT	Minutes-In-Trail
MIT	Miles-In-Trail
NAS	National Airspace System
NEMS	NAS Enterprise Messaging System

SFDPS	System Wide Information Management (SWIM) Flight Data Publication Service
STMP	Special Traffic Management Programs
SVM	Support Vector Machines
TFMS	Traffic Flow Management System
TMA	Traffic Management Advisor
TMI	Traffic Management Initiatives
TRACON	Terminal Radar Approach Control

SUMMARY

One of the missions of the Federal Aviation Administration (FAA) is to maintain the efficiency and the safety of the National Airspace System (NAS). One way to do so is through Traffic Management Initiatives (TMIs). TMIs, such as reroute advisories, are issued by Air Traffic Controllers whenever there is a need to balance demand with capacity in the National Airspace System. Indeed, rerouting flights ensures that aircraft comply with the air traffic flow, remain away from closed airspace, and avoid saturated areas of the airspace and areas of inclement weather. Reroute advisories are defined by their level of urgency i.e. Required, Recommended or For Your Information (FYI). While pilots almost always comply with required reroutes, their decisions to follow recommended reroutes vary. Understanding the efficiency and relevance of recommended reroutes is key to the identification and definition of future reroute options. In addition, since the traffic situation in the national airspace can be forecasted with airlines and Air Traffic Controller (ATC) schedules it would be interesting for airlines and ATC to predict the issuance of volume-related reroute advisories.

Consequently, the objective of this work was two-fold: 1) Assess the relevancy of existing recommended reroutes, and 2) predict the issuance and the type of volume-related reroute advisories.

The first objective has been fulfilled first by acquiring, processing and fusing four months (January – April 2017) of two datasets: System Wide Information Management Flight Data Publication Service (SFDPS) and Traffic Flow Management System (TFMS). The author then assessed the compliance of flight to recommended reroutes based on two approaches: definition of polygon around the reroute and definition of circles around each element of the

reroute. Four compliance metrics have then been developed and implemented to assess the compliance of all flights affected by recommended reroutes between January and April 2017. Results obtained show that very few flights comply with the recommended reroute they are affected by. Further analysis has then been conducted to filter flights according to distances flown and airline types in order to identify trends in the compliance of flights to recommended reroutes.

The second objective has been fulfilled by first fusing traffic data (hourly traffic count per facility) and volume-related reroute advisories extracted from TFMS. Seven Machine Learning algorithms known for their classification abilities have then ben benchmarked on two predictions:

- The prediction of the issuance of volume-related reroute advisories.
- The prediction of the issuance and the type (Required, Recommended, FYI) of volume-related reroute advisories

For both predictions, the best performing technique has been identified in order to have two prediction models.

CHAPTER 1. INTRODUCTION

1.1 The National Airspace System

The National Airspace System (NAS) is comprised of air navigation entities, air traffic controllers, facilities, landing areas, technologies, rules, regulations and procedures. These are needed to manage and ensure the safety of the United States airspace [1]. The US airspace itself is broken down into twenty-one sectors (Figure 1), with each one having precise characteristics in terms of capacity and traffic. Hence, when demand exceeds a sector's capacity, the air traffic in that sector is said to be congested and Air Traffic Controllers have to take actions in order to bring the traffic back to a normal situation.

There is a large variety of entities responsible for the efficiency and safety of NAS operations. They include:

- Air Traffic Control System Command Center (ATCSCC), which manages the NAS in a safe, efficient and cohesive manner [2]
- 21 Air Route Traffic Control Centers (ARTCC), which manage the air traffic of the airspace [2]. Each ARTCC is responsible for one navigation facility of the national airspace. These facilities are presented in Figure 1.
- Terminal Radar Approach Control (TRACON) centers, which control aircraft approaching and leaving (5nm to 50 nm) any airport [2]
- Airport towers which control the flow of aircraft within 5 nm of the airport [2]

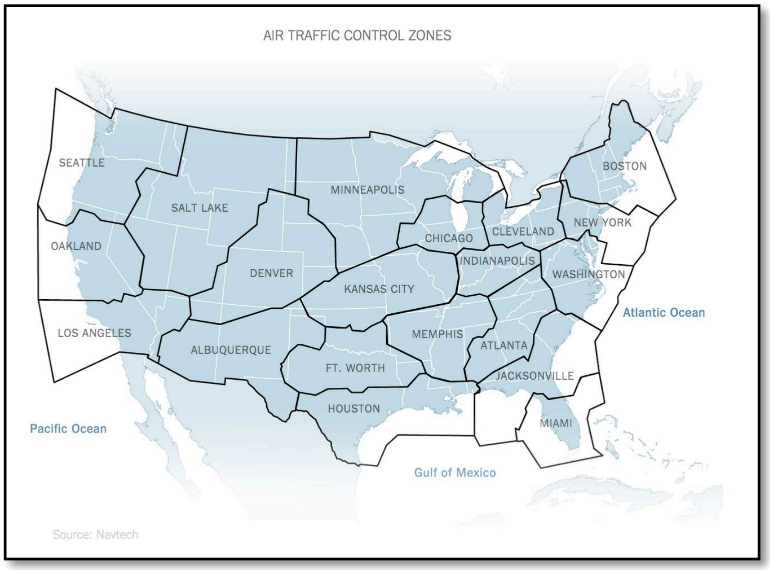

Figure 1: Air Traffic Control Facilities of the NAS [3]

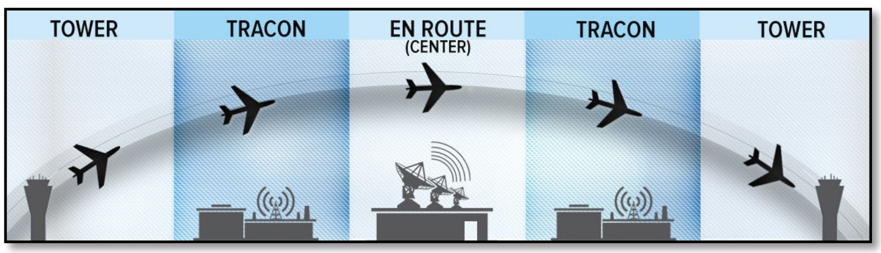

[4]

Figure 2: Overview of facilities controlling aircraft during phases of flight

Figure 2 shows which facility manages aircraft at various stages of their flight. The chronological order from departure to landing (Figure 2) is: Airport Tower, TRACON, ATRCC, TRACON, Airport Tower.

On a daily basis, the National Airspace System is impacted by a variety of events, such as weather, congestion, etc. Congestion, in particular, occurs when the demand for a specific sector exceeds its capacity. One way to mitigate the impact that such events may have on the efficiency and safety of the NAS is through the issuance of Traffic Management Initiatives (TMI).

1.2 Traffic Management Initiatives

Traffic Management Initiatives are in place to "balance demand with capacity either at an airport or in a section of the airspace" [5]. As such, TMIs represent the main means to manage the overall flow of traffic in the National Airspace System. Traffic Management Initiatives are divided into two categories: Airport-Specific (Terminal) Traffic Management Initiatives and En Route Traffic Management Initiatives [6].

1.2.1 Airport Specific (Terminal) Traffic Management Initiatives [7]

These Traffic Management Initiatives are issued to deal with the flow of aircraft arriving at an airport. If the number of aircraft heading to an airport is above the airport's Airport Acceptance Rate (AAR), air traffic managers may issue any of the following Traffic Management Initiatives to slow down air traffic and ensure that the airport's acceptance rate matches or exceeds aircraft demand.

- Ground Delay Programs (GDP) are issued when the "projected traffic demand of an airport is expected to exceed the airport's acceptance rate for a long period of time" [8]
- Ground Stops (GS) are issued when the "projected traffic demand of an airport is expected to exceed the airport's acceptance rate for a short period of time" [8]
- Special Traffic Management Programs (STMP) are issued whenever special events are projected to generate high demand at an airport

1.2.2 En-Route Traffic Management Initiatives [7]

Traffic Management Initiatives may be issued to manage active flights affected by constraints in the National Airspace System. These include:

- Airspace Flow Program (AFP), which identifies constraints in the en-route sector of the National Airspace System and issues a live-time list of all flights affected by the constraint [8]
- Miles-in-Trail (MIT) / Minutes-in-Trail (MINIT) initiatives, which describe the distance in miles or the time in minutes required between two aircraft in a constrained area. [8]
- Traffic Management Advisor (TMA), which schedules aircraft to the active runway threshold of an airport with minimal delay
- Collaborative Trajectory Options Program (CTOP), which automatically affects delay or reroute over constrained area to balance demand and capacity
- Reroutes, which are used to issue new routes to aircraft that need to be diverted from or into a sector of the National Airspace System

In order for each flight operator to understand the new routes issued by Reroute Traffic Management Initiatives, a standardized format is used, as described in the following section.

1.3 National Airspace System Routes

1.3.1 Types of Routes

Routes used by flight operators can be classified into three groups [6]:

- *Preferred Routes:* These are common routes between two airports that pilots regularly use in the absence of constraints in the National Airspace System. These were designed to increase the National Airspace System's efficiency and capacity
- *Playbook Routes:* These are routes that have been created and pre-validated to fit particular circumstances. They are used when preferred routes are not available because of any incident
- *Coded Departure Routes:* These are coded air traffic routes that are used to reduce the amount of information transmitted and making communication between Air Traffic Controllers and flight operators more efficient

1.3.2 Format of routes

Flight routes follow an internationally standardized format to ensure uniformity across the world. Their format is a sequence of elements that belong to a catalogue of points and routes accredited by the aviation administration they belong to. Below is a sample route extracted from a flight plan:

$$KEWR..ELVAE..COL..WHITE.J209.SBY..KEMPR..ILM.AR21.CRANS.FISEL6.KFLL$$

Flight routes between two airports always start and end with the origin airport and the destination airport: in the example provided above, KEWR (Newark Liberty International Airport) is the origin and KFLL (Fort Lauderdale–Hollywood International Airport) is the destination. This route is represented in Figure 3.

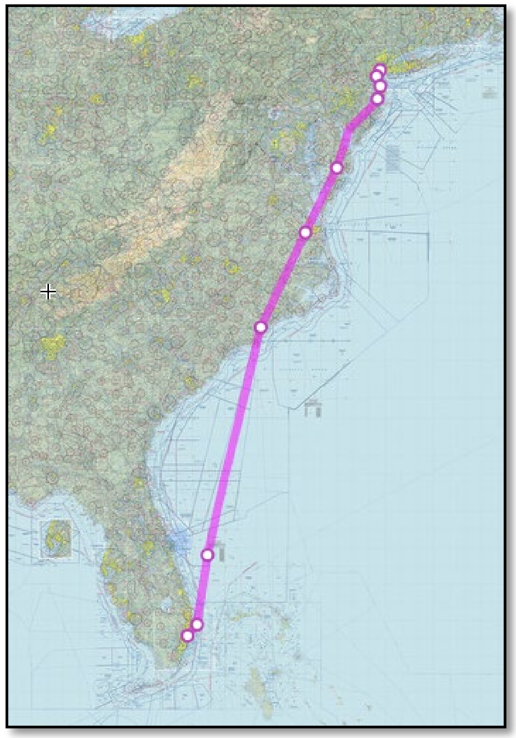 [9]

Figure 3: Map visualization of a route from KEWR to KFLL

ELVAE, WHITE, KEMPR, CRANS are FAA Fix Waypoints. These waypoints are geographical points on the Earth's surface.

COL, SBY, ILM are Navaids, which are physical devices on the ground that transmit radio signals that aircraft can detect and follow.

Finally, *J209* and *AR21* are air route names. They are alphanumeric codes that define corridors connecting specified locations to each other at specified altitudes.

Based on the definition and explanation of the National Airspace System's routes and their formats, it is possible to understand and analyse Reroutes Traffic Management Initiatives. As aforementioned, these Reroute advisories are issued by Air Traffic

Controllers. However, the Federal Aviation Administration currently does not have any information as to the relevancy of these Reroute advisories. Consequently, the present research proposes to focus specifically on Reroute Traffic Management Initiatives with the objective to provide FAA analysts with increased awareness as to the relevancy of reroutes.

1.4 Reroute Advisories

As aforementioned, reroute advisories are Traffic Management Initiatives (TMI) that are issued when an Air Route Traffic Control Center (ARTCC) identifies constraints in the National Airspace System. Under such TMIs new routes are assigned to affected flights. Rerouting flights ensures that aircraft comply with the air traffic flow, remain away from closed airspace such as those for military use, avoid overcrowded areas of the airspace, and avoid areas of inclement weather. The following sections discuss the whole process that lead Air traffic Controllers to issue and define Reroute Traffic Management Initiatives.

1.4.1 Rerouting Process

An area in the National Airspace System can be constrained for various reasons, including, but not limited to, inclement weather conditions or aircraft congestion. Whenever an area is constrained, traffic management personnel locate the constraint, assess which airport(s) and route(s) are affected, and evaluate the seriousness of the constraint and its duration. Once this information has been gathered, an ARTCC can decide to issue Traffic Management Initiatives such as reroutes to address the constraint. Reroute advisories specify the constrained area, the effective period of the advisory, the nature of the incident, the probability of extension and the new routes for affected flights. Once flight operators receive a reroute advisory, they then have to either submit a flight plan amendment or submit an alternative route and check with Air Traffic Controllers if it is accepted.

1.4.2 Rerouting levels of urgency

Reroute advisories issued by ARTCCs contain several characteristics about the incident, the new routes, and their level of urgency. There are three different levels of urgency for reroute advisories, with each one having different requirements on flight operator's compliance [6]:

- *Required Reroutes:* These routes are required to be followed by all aircraft captured in the scope of the reroute
- *Recommended Reroutes:* Air Traffic Controllers recommend flight operators to use these routes, but do not require them to use them
- *For Your Information (FYI) Reroutes:* Air Traffic Controllers issue these reroutes to let pilots know that these routes are available

Figure 4 shows the distribution of reroute advisories per urgency level and for different causes for one day (April 21, 2017).

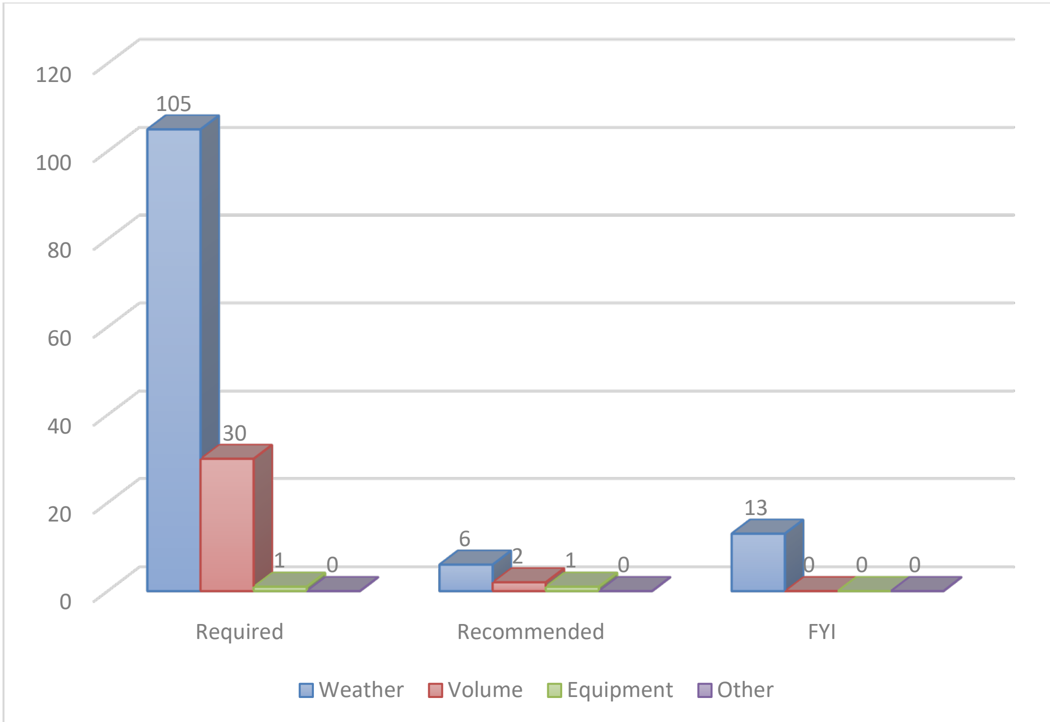

Figure 4: Reroute Advisory Statistics (April 21, 2017)

In particular, it shows that even if a large majority of reroute advisories are required, Air Traffic Controllers still issue around twenty recommended and FYI reroute advisories per day, which is not negligible.

1.5 Research Scope & Objective

1.5.1 The relevancy of recommended reroutes

As mentioned previously, reroute advisories are defined by their level urgency that can be either required, recommended or just informative. According to Federal Air Regulation (FAR) §91.123 [10], flight operators are allowed to refuse a specific route as long as they submit an alternative one that is validated by Air Traffic Controllers. The Federal Aviation Administration (FAA) analysts are interested in analyzing how often pilots follow recommended reroutes in order to assess their relevancy, and eventually define reroutes that are more likely to be followed. Doing so would eventually lead to reductions in flight delays and flight durations.

1.5.2 Prediction of reroute advisories

As discussed, Traffic Management Initiatives such as reroutes can be issued if constraints appear in the NAS. These constraints are mostly related to the following reasons: weather conditions, volume or equipment. Reroute advisories due to volume constraint are issued when the traffic in an ARTCC exceeds the capacity of the facility. Since air traffic flow can be predicted with airlines schedule, so can volume-related reroutes. Being able to predict issuances of reroutes due to volume and eventually the type of the reroute advisory will assist flight operators and traffic flow management personnel to plan routes more efficiently to avoid further unnecessary reroutes. Weather-related and equipment-related reroutes will not be included in the prediction model for different reasons. Predict reroutes due to weather conditions would make sense to many airlines because it is possible to have weather forecasts. However, building the predictive model requires a large amount of training data that is hard to collect because reroute advisories do not provide a detailed statement, with location and duration for example, about the weather conditions causing the reroute. Equipment issues, on the other hand, are very hard to forecast and it is therefore very complex to build a reliable prediction of the issuance of equipment-related reroute advisories.

1.5.3 Research Objectives

The objective of this research is thus two-fold: 1) Assess the relevancy of existing recommended reroute advisories issued by Air Traffic Controllers, and 2) predict the issuance of reroute advisories due to volume constraints.

In order to fulfill these objectives, access and analyze relevant datasets that are currently used by the Federal Aviation Administration (FAA) is one path. This data, as further discussed in the following chapters, has the characteristics of what is commonly referred to as "Big Data". Hence, the following chapter introduces Big Data problematics and their importance for the Federal Aviation Administration. It also discusses past efforts as they relate to the analysis of flight reroutes to help define the research gaps to be addressed by this research.

CHAPTER 2. BACKGROUND

The FAA generates, receives, stores and utilizes very large amounts of data coming from diverse sources across industry and government, including airlines, airports and Air Route Traffic Control Centers (ARTCC). Each flight, for example, generates approximately half a terabyte of data which can then be used by airlines and manufacturer alike to get insights about the performance, reliability and maintenance of aircraft [11]. This data, commonly referred to as Big Data, is of critical value to aviation stakeholders. If acquired, stored, processed and analyzed properly, Big Data may lead to the following benefits:

- Identified patterns for development of new concepts and methods
- Real time analysis and prediction of air traffic operations
- Improved flight safety and efficiency in the National Airspace

2.1 Big Data

Big Data is referred as "large and complex massive amounts of datasets that it becomes difficult to process and analyze using traditional data processing technology." [12]. It is characterized by the four V's:

- Volume: Refers to the size and the amount of data collected. Big Data involves datasets ranging from terabytes (10^{12}) to zettabytes (10^{21})
- Velocity: Refers to the frequency and speed at which data is streamed and collected
- Variety: Refers to the various data formats

- Veracity: Large datasets might contain significant amount of imprecise and uncertain data that needs to be extracted before analysis

2.1.1 Federal Aviation Administration and Big Data

Big Data is used by the FAA to ensure that commercial and general aviation over the United States is the safest in the world. Hence, to efficiently manage the air traffic comprised of approximately 42,000 daily flights [4], the FAA needs methods for analyzing Big Data. To store, process and analyze Big Data, FAA analysts and researchers use advanced data analytics methods such as Data Fusion, Data Parsing and Machine Learning.

2.2 Machine Learning

Machine Learning is now considered as a complete subset of computer science, and has been defined by Arthur Samuel, a Machine Learning pioneer, in 1959 as:

"The field of study that gives computers the ability to learn without being explicitly programmed." [13]

Machine Learning is widely used in various industries because of its strong abilities in data analytics. Examples of machine learning applications in the industry include analysis of medical images [14], optimization of car users' routes by predicting traffic [15], prediction of Uber estimated time of arrival and prices [16] or reduction of fraudulent credit card transactions [17], to name just a few.

The evolution of Machine Learning techniques over the past 50 years is due to intertwined factors: improvements in computing/processing power, decreases in the cost of acquiring and storing data, the need to analyze even larger datasets, etc.

The benefits brought on by Machine Learning is dependent on the quality (and sometime quantity) of the data acquired as well as the proper selection and training of machine learning techniques. Indeed, each Machine Learning algorithm relies on various assumptions that need to be understood and parameters that need to be tuned. Having a clear understanding of these assumptions is thus critical to the correct application of any machine learning algorithm.

Machine Learning algorithms are divided into three subsets that have different assumptions and different fields of applications:

- *Supervised Learning:* These algorithms predict a target value using other values of the dataset. In supervised learning, the algorithm attempts to discover and model relationships between the target value and other values. Supervised Learning is very efficient to solve problem dealing with two types of data analytics: classification and numeric predictions [18].
- *Unsupervised Learning:* These algorithms study datasets on their own to identify patterns, determine correlations and relationships. Models used in unsupervised learning are descriptive models. With this kind of algorithm, as opposed to supervised algorithms, there is no labelled target features. These algorithms are known for their ability on clustering or pattern recognition problems [19].

- *Reinforcement Learning:* These algorithms are learning what to do and how to interact with their environment so as to optimize a numerical reward signal. Reinforcement learning is different from Supervised Learning because the learner is not told what to do, but instead has to try each possible action and figure out which one yields the highest reward. The learner also needs to be aware that actions may have an impact not only on the immediate reward but also on all following rewards. Reinforcement learning main features and differences with other Machine Learning types are therefore the trial-and-error search and the delayed reward [20].

The following sections review past studies that support the objective of this research and discuss existing research gaps.

2.3 Review of prior research related to rerouting advisories

Few studies have been conducted to assess the relevancy of recommended reroutes. The most relevant ones are discussed below.

2.3.1 Rerouting algorithm

Most of the prior research related to rerouting advisories focused on developing algorithms that created reroutes and improved or created new processes to affect reroutes.

- *"ARTCC Initiated Rerouting", 2006* [21]

This effort analyzed the process of reroute planning and execution. The ultimate goal was to develop a new rerouting process for ARTCC in order to increase the common situational awareness for potential local reroutes. This would allow NAS users to submit reroute alternatives and increase the automation support for ARTCC to identify, assess and execute a reroute. Even though the experimental results validated the new process, it is interesting to note that this work was completed before the implementation of the Traffic Flow Management System (TFMS) in 2008. It means that the FAA was also working on improving their traffic management system and the old model this work is based on is now obsolete.

- *"Robust Air Traffic Control Using Ground Delays and Rerouting of flights", 2009* [22]

This effort assessed the efficiency of ground delays and rerouting advisories for three scenarios that may affect the performance of the National Airspace System: loss of ARTCC, loss of a link between two ARTCC and isolation of an ARTCC over a period of time. The main metric used to assess the efficiency was the number of aircraft that needed to be diverted in order to restore the performance of the National Airspace System. This work involved developing two models comprised of seven and twenty ARTCCs. In both cases, it appeared that ground delays and rerouting were efficient to mitigate the effect of an ARTCC going down in the NAS. However, from a computational point of view, a ground delay-based optimization approach was significantly less complex than a rerouting optimization approach, which generally resulted in nonlinear programming problem.

- *"Pilot Convective Weather Decision Making in En Route Airspace", 2012* [23]

This effort examined the strategic aspects of pilots' behaviors during a weather avoidance process in the tactical time frame (0 – 2 hours from the incident). The goal of this study was to implement an algorithm to increase the automation level of a rerouting process involving pilots and Air Traffic Controllers. This work involved asking eighteen transport pilots to participate in lab studies where they were presented with a weather encounter scenario in an en-route environment. In particular, these pilots were asked to modify the planned trajectory in the event that they found it unsafe given the weather forecast. Results of these simulations showed that pilots were more willing to trade safety for flight efficiency even if it implied not respecting FAA guidelines on separation assurance. The scope of this study can be extended to non-weather related reroutes.

2.3.2 Traffic Management Initiatives Statistics

"Aggregate Statistics of National Traffic Management Initiatives", 2010 [24]

This effort aggregated and analyzed data from the National Traffic Management Log to provide a set of statistics on the implementation of Traffic Management Initiatives. Extracted statistics on reroutes focused on ranking points such as airports, waypoints and navaids according to the number of reroutes they are affected by. Results showed that the five points most affected by reroutes were airports (Denver, Newark, Dallas, JFK and La Guardia). Denver and New York area airports (EWR and JFK) were affected by approximately two reroute advisories per day for instance. This effort also consisted in developing a visualization tool aimed at increasing both understanding and situational awareness of Traffic Management Initiatives. While this effort statistically described Traffic Management Initiatives, it did not assess the efficiency or the relevance of Traffic Management Initiatives.

2.3.3 Research on adherence of flights to routes

In order to assess the relevancy of recommended flight reroutes, it is interesting to have an overview of research about the adherence of flights to routes. Two studies are presented below.

"Determination of Lateral and Vertical Adherence to Route", 2013 [25]

This study developed a new algorithm to evaluate the adherence of an aircraft to a route by calculating the lateral and vertical deviation. These deviations are calculated based on the definition of thresholds. For the lateral deviation, the inner threshold and the outer threshold are defined such that if the lateral distance from the plane's track position to the route is inferior to the inner threshold, the aircraft is assumed to be in adherence. If the lateral distance is superior to the outer threshold, the aircraft is assumed to be out of adherence. If the lateral distance is between the two thresholds, the aircraft must show the willingness of returning to its route to be assumed in adherence. The authors developed an algorithm to assess this intent. The algorithm is based on the comparison of the direction of the plane and the direction of the route. The end of the study focuses on calibrating all parameters and find the optimum values for each threshold based on a truth reference dataset of 100 flights semi-randomly selected from a set of 2,234 flights.

This effort provides an interesting approach on the comparison of a flight trajectory and an air route. However, while the general idea and the description of the metric are very detailed, the testing phase and in particular the truth reference dataset are not as well described. It would be interesting to know how much the data format influences the implementation of the metric and what was the content of the dataset they collected from the Chicago Air Route Traffic Control Center.

"Evaluation of new enroute performance measures for air navigation service providers", 2017 [26]

This effort focuses on designing and evaluating new metrics to assess the performance of Air Navigation. It focuses particularly on the lateral deviation between flight trajectories and flight plans or routes. The author's approach in this study relies on building horizontal area between the flight trajectory and a route in order to determine the lateral proximity. The horizontal area is built as the sum of areas of the polygons between each intersection between the flight trajectory and the route. The lateral deviation is then calculated as the horizontal area divided by the length of the route.

The approach developed in this effort is different from the one in the previous effort and it is hard to say that one is more efficient than the other. They are both interesting and certainly helpful for the author if there is a need to compare flight trajectory and recommended reroutes.

2.4 Summary of prior research and research gaps

The review of prior research conducted on the topic of Traffic Management Initiatives, and rerouting in particular, highlights some limitations and gaps. First, most of the research conducted so far focused on developing and/or improving the traffic management rerouting process. This has been achieved by either creating a new reroute creation algorithm or determining the most efficient reroute using an optimization algorithm. Some studies also focused on assessing the adherence of flights to routes by developing adherence metrics but none of them focused precisely on assessing the compliance of flights to reroutes. In addition, while the adherence of flights to routes can be assessed with various metrics, the implementation of those metrics is strongly related to the data that has been collected. Consequently, the present research aims to address this limitation by identifying relevant and modern datasets providing reroute advisories and traffic data, developing compliance metrics easily implementable with the analyzed data and finally extracting statistics and trends that may lead to the assessment of the relevancy of reroute advisories.

Second, most of previous efforts have only focused on weather-related reroutes, ignoring other causes such as traffic volume or equipment issues. However, traffic volume constraints are also an important reason for reroutes and should not be ignored. This research aims to address this limitation by analyzing the relevancy of all recommended reroute advisories.

Third, previous efforts have not distinguished the urgency level of reroutes: required, recommended or FYI. This may have occurred because previous efforts did not have access to the data needed for such a study. This research aims to address this limitation by extracting reroute advisories from Traffic Flow Management System (TFMS) datasets and analyzing recommended reroutes.

Finally, no research has been conducted on the prediction of the issuance of reroute advisories. This research aims to fill this gap by developing a prediction model to predict the issuance of volume-related reroute advisories.

CHAPTER 3. PROBLEM DEFINITION

3.1 Assessment of the relevancy of recommended reroutes

As discussed in Chapter 2, no work has been conducted that assesses the relevancy of recommended reroutes. There is thus, a need to address this gap. Recommended reroutes, as opposed to Required or FYI reroutes, are of particular interest to the FAA because 1) the decision to follow them is left to the pilots, and 2) it concerns a substantial amount of traffic. Hence, the first question this research seeks to answer focuses on assessing the relevancy of recommended reroutes.

Research Question 1: *How can the relevancy of recommended reroutes be best captured?*

Lack of access to relevant and comprehensive traffic information and data, flight information, and weather reports has limited the ability of researchers to assess the efficiency and the relevance of recommended reroutes. Many datasets identified by the author and made available by the FAA have the potential to help address Research Question 1. Datasets such as the System Wide Information (SWIM) Flight Data Publication Service (SFDPS) and Traffic Flow Management System (TFMS), for example, provide relevant information for assessing the relevancy of recommended reroutes.

System Wide Information Management (SWIM) Flight Data Publication Service (SFDPS)

The System Wide Information Management Flight Data Publication Service (SFDPS) dataset provides individual flight information about en-route aircraft to National Airspace

System stakeholders. It allows stakeholders to receive and process real-time data for informational, analytics, research or any other purpose related to air traffic over the NAS. SFDPS gathers Service-Oriented (SOA) message patterns in order to publish data from the En-Route Automation Modernization (ERAM) system. ERAM data is issued through the Host Air Traffic Management (ATM) Data Distribution System (HADDS), which is one element of the En-Route Data Distribution System (EDDS). Each one of the 21 Air Route Traffic Control Centers (ARTCC) hosts these system [27].

SFDPS messages are divided into three subsets in FIXM format:

- *SFDPS Derived Messages*: These messages are created by SFDPS to provide answers to stakeholder requests or to provide system status information
- *Reconstitution Messages*: Reconstitution Messages are received from the Host Air Traffic Management (ATM) Data Distribution System (HADDS) and are stored in the Database Record Transfer (DRBT).
- *Flight Data Messages*: These messages include any relevant information about each individual flight in the National Airspace System such as flight plans, track data for active flights, arrival and departure information, etc.

Each of the aforementioned message groups consists of different messages containing different information. Within the SFDPS dataset, the *Flight Data Messages* subset appears to be the most interesting one for the purpose of this research. Indeed, it contains data specific to individual flights such as tracking positions, altitude, speed and flight plan that can be used to help assess the relevancy of recommended reroute advisories.

Traffic Flow Management System (TFMS)

The Traffic Flow Management System (TFMS) "predicts, on national and local scales, traffic surges, gaps, and volume based on current and anticipated airborne aircraft. Traffic management specialists evaluate the projected flow of traffic into airports, sectors, and fixes, and then implement the least restrictive action necessary to ensure that traffic demand does not exceed system capacity" [28]. TFMS also provides Aircraft Situation Display (ASDI) data such as aircraft scheduling, routing and positional information. TFMS is comprised of two subsets: TFMS Flight and TFMS Flow [29].

- TFMS Flight provides data related to flights being managed by TFMS and is made up of the following elements:
  - Flight Plan Data and potential updates and amendments
  - Departure & Arrival time notifications
  - Flight cancellations
  - Boundary crossings
  - Track position records
  - Flight management
  - NAS common situational model data
  - Flight Table Manager deltas
- TFMS Flow provides the definition of Traffic Management Initiatives, changes to the definitions, and their cancellations. TFMS Flow is comprised of the following messages:
  - Traffic Management Initiative definitions

- Ground Delay Program / Unified Delay Program
- Airspace Flow Program
- Collaborative Trajectory Options Program
- Flow Constrained Area / Flow Evaluation Area definitions
- ATCSCC advisories
- Restrictions
- Airport runway configuration and rates
- Airport deicing status
- Route availability planning tool timeline data

Once datasets are identified and acquired, there is a need to understand their content and identify how information should be combined to provide analysts with the big picture. Data fusion is one approach that can be used to help address this challenge. As discussed by Sorber, Van Barel and Lathauwer [30], data fusion is "the process of integrating and analyzing data from multiple sources in order to develop insights that are deeper and more accurate than those resulting from a single source of data".

Thus, with the identification of suitable datasets and Data Fusion, the following hypothesis can be made to answer *Research Question 1*:

Research Question 1: *How can the relevancy of recommended reroutes be best captured?*

Hypothesis 1: *If Data Fusion is used to analyze recommended reroute messages and flight data, then metrics may be defined and implemented to assess the relevancy of recommended reroutes.*

3.2 Prediction of the issuance of reroute advisories due to volume constraints

Reroute advisories may be issued as Traffic Management Initiatives by air traffic controllers to address certain constraints in the National Airspace System. These reroute advisories can be issued for diverse reasons but those three are the most common: weather conditions, volume constraints, equipment issues. As stated Section 1.5.2, the nature of the data available and the complexity to forecast equipment issues led the author to only focus on the prediction of volume-related reroute advisories. This leads to the second research question this work is trying to answer:

Research Question 2: *How can the issuance of a volume-related reroute advisory be accurately predicted?*

For the scope of this research question, there is no valid reason to analyze exclusively recommended reroutes whereas it was meaningful for the Research Question 1 because the vast majority of flights complies with required reroutes. Indeed, it might be even more interesting for Air Traffic Controllers but also for airlines to be able to predict the issuance of required reroutes than recommended reroutes because of their mandatory nature. Therefore, for the scope of this question, no reroute advisory type (RQD, RMD, FYI) will be filtered out and the final predictive model will be as comprehensive as possible.

Volume-related reroutes are issued by Air Traffic Controllers when an area of the airspace is expected to be congested or when the air traffic exceeds the capacity of a sector. Air Traffic Controllers control the air traffic with the Flow Evaluation Area (FEA) and Flow Constrained Area (FCA) tools. The FEA consists in separating portions of the airspace with lines or polygons and then compute the amount of traffic across those portions. When the amount of traffic is considered to be a potential volume issue in an area, the FEA becomes an FCA and Traffic Management Initiatives such as reroute advisories can be issued to avoid the volume issue.

In order to predict the issuance of a reroute advisory due to volume constraint, there is first a need to gather large amounts of historical data about the air traffic in the national airspace and reroute data. Reroute data is obtained from TFMS dataset as explained previously. Air traffic data such as traffic counts for ARTCC will be obtained from the FAA.

Predicting whether a volume-related reroute advisory will be issued can be represented as a classification problem. Supervised Learning algorithms in general and Support Vector Machines (SVM), Neural Networks or Decision Trees, in particular, are known to be suitable methods for solving classification problems. Many benchmarking studies have been conducted to compare machine learning algorithms for classification problems. However, the performance of a machine learning algorithm is strongly related to the input dataset, therefore results of one study might not be reliable and true for a different dataset. Thus, additional efforts need to focus on testing and comparing different machine learning algorithm to predict the issuance of reroute advisories due to volume constraints. The machine learning algorithms investigated include: Naïve Bays Classifier, Support Vector Machines, Decision Trees, Boosting Trees, Bagging Ensembles, Random Forest and Nearest Neighbor.

Consequently, the following hypothesis can be formulated:

Research Question 2: *How can the issuance of a volume-related reroute advisory be accurately predicted?*

Hypothesis 2: *If traffic data is fused with reroute data (TFMS), and supervised machine learning algorithms are used to develop prediction models, then it will be possible to find the algorithm that best predicts the issuance of volume-related reroute advisories.*

CHAPTER 4. PROPOSED APPROACH

A six-step methodology is proposed to test the aforementioned hypotheses and answer the research questions enunciated in the previous chapter. The following sections discuss each step in detail.

4.1 Step #1: Data identification and acquisition

4.1.1 Datasets

As mentioned in the previous chapter, two main datasets are considered for the scope of this research: System Wide Information Management (SWIM) Flight Data Publication Service (SFDPS) and Traffic Flow Management System (TFMS). These datasets are obtained from the FAA's Computing Analytics and Shared Services Integrated Environment (CASSIE) and include data from January 2017 to April 2017. In order to fully understand reroute advisories extracted from TFMS files, two more datasets storing coordinates of geographic points used to describe reroutes are needed. These datasets can be found online [31] and are provided by the National Airspace System Resource (NASR) System.

The SFDPS and TFMS datasets are in the Flight Information Exchange Model (FIXM) format, which captures flight and flow information that is globally standardized [32].

System Wide Information Management Flight Data Publication Service (SFDPS)

SFDPS is comprised of Flight Data messages containing many different message types such as Flight Plan, Flight Amendment Information, Cancellation Information or Hold Information for example, but the most relevant one for the scope of this research is Track Information (TH) messages [27]:

> *Track Information (TH_FIXM):* These messages provide flight track data such as aircraft position, altitude, and speed every twelve seconds for active flights

All TH_FIXM messages have the same structure and are providing the same information for each flight. Within each message, the following fields are extracted because they provide relevant information for this research [27]:

- *propMessageType*: specifies the message type received from the HADDS
- *propFlightId*: specifies flight numbers
- *propOrigin:* specifies the origin of flights
- *propDestination*: specifies the destination of flights
- *propSentTime*: specifies the time at which the message was sent from SFDPS to the NAS Enterprise Messaging System (NEMS)
- *arrivalTime*: specifies the expected arrival time of flights
- *departureTime*: specifies the departure time of flights
- *flightState*: specifies the status of the flight. This can be either Active, Cancelled, Dropped, Landed or Proposed

- *trackPosition_23d*: specifies the real-time position of the flight as a latitude/longitude pair
- *reportedAlt_54a*: specifies the reported altitude of the aircraft in hundreds of feet

Traffic Flow Management System (TFMS)

As mentioned previously, the TFMS Flow dataset contains advisories that were issued as part of Traffic Management Initiatives. For the purpose of this research, General Advisory (GADV) messages are extracted and used because they contain the recommended reroute advisories. The following fields contained in GADV messages provide relevant information regarding recommended reroutes:

- *fcm:advisoryNumber*: specifies the advisory ID number
- *fce:startTime*: specifies the start time of the advisory
- *fce:endTime*: specifies the end time of the advisory
- *fcm:advisoryTitle*: coded sentence that summarizes the advisory
- *fcm:advisoryText*: this contains extensive information on the advisory including the constrained area, the reason, the probability of extension, the new routes, and any relevant remark.

Routes defined in the *'avisoryText'* field can appear in two different formats. The first one, which is used mostly for required reroutes, defines the complete reroute including the original airport, the destination airport and the route between these two airports: fixes, Navaids and airways number. Below is an example of reroutes issued in the first format.

ROUTES		
ORIG	DEST	ROUTE
KMSP	KJFK	DLL HASTE DAFLU J70 LVZ LENDY6
ATL	DTW	VXV J91 HNN DJB GEMNI4
LAS	ATL	INW J86 ELP ABI J4 MEI DUUCK PRICI RAGGZ1

Figure 5: Example of reroutes defined in the first format with origin, destination and route

The second possible format, mostly used for recommended reroutes and which is thus the one that is analyzed the most in this work, defines route origin segments from origin airports and route destination segments for destination airports. Below is an example of routes issued in the second format.

ROUTES	
FROM	
ORIG	ROUTE ORIGIN SEGMENTS
EWR	DIXIE PREPI UNYAD OWENZ POPPN OHRYN BEHHR WEBBB HOBOH PAEPR M201 HANRI
FLL FXE	ZAPPA PERMT AR16 EMCEE M201 PAEPR
MIA TMB	VALLY PERMT AR16 EMCEE M201 PAEPR
PBI BCT SUA	PBI A699 PERMT AR16 EMCEE M201 PAEPR
TO	
DEST	ROUTE DESTINATION SEGMENTS
BCT	HANRI M201 JENKS AR19 AYBID CAYSL4
CDW MMU	PAEPR HOBOH SILLY STINK YAALE YETTI MOUGH DONAA OWENZ CYN GXU RBV V249 METRO
FLL	HANRI M201 BAHAA AR21 CRANS FISEL7

Figure 6: Example of reroutes defined in the second format with origin, origin segments and destination, destination segments

With this format, the full reroute is obtained by compiling the route origin segment and the route destination segment. This format needs some deeper analysis though because it often occurs that the same airport appears in the origin and destination sections and therefore defines a non-sensical airport with the same airport as origin and destination. Moreover, there are sometimes situations where a reroute is defined between two close airport (BOS and JFK for example) and provides a very long and non-sensical route such as the one presented in Figure 5.

Figure 7: Example of non-sensical reroute defined with second reroute format

Thus, for reroutes defined with the second format it is necessary to filter out the non-sensical reroutes, those between two identical airports or between two very close airports, and process only the subset with meaningful reroutes.

As explained in Chapter 1, routes provided in the '*advisoryText'* field are made of fixes, navaids and airways identifiers. In order to visualize these routes and compare them to flight trajectories it is necessary to extract coordinates of each element of the reroute.

Fix/Reporting Point/Waypoint dataset

This dataset is provided by the National Airspace System Resource (NASR) System [31] and is updated every 28 days. It lists all FIX waypoints and provide for each of them record identifier, ICAO region code, latitude/longitude of the fix, military or civil fix and any relevant information describing the fix. For the scope of this research, only the record identifier and the geographical coordinates are extracted.

Navigation Aids Dataset

This dataset is also provided by the National Airspace System Resource (NASR) System [31] and updated every 28 days. It lists all Navigational Aids (Navaids) waypoint and provide for each of them the record identifier, the Navaid facility type (VORTAC, VOR/DME, FAN MARKER, MARINE NDB, etc…), its latitude/longitude and any relevant information describing the Navaid. For the scope of this research, only the record identifier and the geographical coordinates are extracted.

4.2 Step #2: Data Processing

4.2.1 SFDPS and TFMS data processing

As mentioned in section 4.1.1, both SFDPS and TFMS datasets are in Flight Information Exchange Model (FIXM) format. While this format is appropriate for storing and sharing large amounts of data, it is not suitable for data analytics purposes. Thus, for analytical purposes, there is a need to parse SFDPS and TFMS into a much more usable and appropriate format such as Comma Separated Values (CSV). The main advantages of the CSV format over FIXM are its compatibility with data analysis techniques applied with Python. More precisely, many existing Python modules such as Pandas or the CSV module can be leveraged to facilitate the loading and analysis of data in a more efficient manner than if the data was left in the FIXM format. SFDPS and TFMS datasets are stored as hourly files by the FAA and are composed of all messages generated within the hour. Furthermore, SFDPS and TFMS datasets have schemas which dictate the dataset's structure. Using schemas is critical to make sure that all required fields are extracted in their correct formats. Schemas are stored as XML Schema Definition (xsd) files and can be found online or directly from Python command. A Python parser developed by Mangortey et al. [33] for TFMS messages has been updated to achieve the objectives of this research. A parser for SFDPS messages has been developed and is based on a process similar to the one implemented for the TFMS parser. The process is as follows:

1. Since the data is stored hourly and made of all messages generated within the hour, it is important to enclose each file with a header and footer such as <root> and <\root> to be able to distinguish between the beginning and end of the file.
2. The schema location is stipulated at the beginning of each message. It then needs to be extracted from the xsd file.
3. The FIXM file is parsed using the ElementTree Python module [34].
4. The field names and their corresponding values are saved to a Python list and the index of the opening tag attached to each message.
5. Using the index of the opening tag, the developed script parses through the list, identifies for each message the useful fields and values and extracts them.
6. The extracted messages are then appended as lists into a final list.
7. The final list is then converted into a CSV file using the csv.writer function of the csv Python module.
8. Finally, each CSV file is saved for each hour appropriately.

4.2.2 Processing on updated reroute advisories

As stated in Section 1.4, reroute advisories are issued whenever an ARTCC identifies constraint(s) and assigns new routes to affected flights. However, constraints such as bad weather conditions, aircraft congestion or equipment issues may change, leading to updates to the scope of reroute advisories. These updates are issued as new reroute messages with new advisory numbers. It is stated in the remarks field of the advisory if it is an update to a previous reroute advisory.

A major challenge is that reroute advisories are still generated after being updated, which makes them invalid. Therefore, there is a need to exclude reroutes advisories after they have been updated. To do so, the end time of the initial reroute advisory will be set as the start time of the updating advisory. Doing so helps ensure that the effective periods of both initial and updating advisories do not overlap and that they are treated as two separates reroute advisories.

Below is an example of an updating advisory with changes appearing in red in the Table 2. In this case the end time of the updated advisory (#71) becomes the start time of the updating advisory (#117).

```
<ds:fiOutput>
     <fi:fiMessage sensitivity="A" sourceFacility="TSS" sourceTimeStamp="2017-04-20T18:37:24Z"
msgType="GADV" refresh="RFRS">
          <fi:generalAdvisory>
             <fcm:advisoryNumber>0117</fcm:advisoryNumber>
             <fcm:origin>ATCSCC</fcm:origin>
             <fcm:dateSent>2017-04-20T18:37:24Z</fcm:dateSent>
             <fcm:facilities>ZBW/ZDC/ZJX/ZMA/ZNY/ZWY</fcm:facilities>
             <fcm:effectivePeriod>
                <fce:startTime>2017-04-20T18:15:00Z</fce:startTime>
                <fce:endTime>2017-04-21T03:00:00Z</fce:endTime>
             </fcm:effectivePeriod>
             <fcm:advisoryTitle>ATCSCC ADVZY 117 DCC 04/20/17 ROUTE RMD</fcm:advisoryTitle>
             <fcm:advisoryText>NAME AZEZUMODIFIEDVIAM202
REMARKS REPLACES ADVZY 071 ZNY ADVISES M201 IS CLOSED DUE TO
      WEATHER A/C LOOKING TO FILE M202 NEED HF AND CAN EXPECT AT
      OR BELOW FL280 FOR ZNY WX
```


Figure 8: Example of an updating advisory

Table 1: Effective period for both advisories 71 and 117 before processing

Advisory Number	Start Time	End Time
0071	2017-04-20T16:00:00Z	2017-04-21T03:00:00Z
0117	2017-04-20T18:15:00Z	2017-04-21T03:00:00Z

Table 2: Effective period for both advisories 71 and 117 after processing

Advisory Number	Start Time	End Time
0071	2017-04-20T16:00:00Z	2017-04-20T18:15:00Z
0117	2017-04-20T18:15:00Z	2017-04-21T03:00:00Z

4.3 Step #3: Data Fusion

The data fusion process involves understanding how the different datasets and their features are related to each other. This involves identifying common fields in order to fuse the data. The first common feature between the two main datasets (SFDPS, TFMS) is time.

The date, time and effective period of recommended reroute advisories are extracted from TFMS, while flight departure and arrival times are extracted from SFDPS.

Another common field between the two datasets is the location. Recommended reroute advisories from TFMS state the affected airport(s) and/or area(s) of the airspace. The origin and destination airport(s) of flights are extracted from SFDPS.

4.3.1 Data Fusion to assess the relevancy of recommended reroutes

The steps taken to fuse the SFDPS and TFMS datasets are as follows:

1. From TFMS recommended reroute advisories, extract the affected airports (departure and arrival), the effective period of the advisory and the suggested routes.
2. Convert each recommended reroute from a sequence of waypoints and Navaids to a sequence of GPS positions corresponding to the actual positions of all waypoints and Navaids.
3. From SFDPS, extract messages of all flights flying from and to the airports affected by the reroute advisory.
4. Within these flights, only keep the one flying during the effective period of the advisory.
5. Within each affected flight, only keep and order TH_FIXM messages by generation time.
6. For TH messages, create a list of the path taken by the flight from origin to destination using flight coordinates.

4.4 Step #4: Data Analysis and Results

To analyze the relevancy of recommended reroute advisories, actual flight paths can be compared to recommended reroute advisories. This can be done in two ways:

1. Flight plan approach: compare flight plans (FH, AH, HU messages) to recommended reroutes.
2. Tracking Flights approach: Track flights using flight coordinates and compare to the path of recommended reroutes.

For the scope of this research, the second approach is the only one that is implemented. The justification for this is that, based on discussions with FAA data analysts, it appeared that this second approach is the most precise and reliable as pilots do not always keep their flight plan updated.

4.4.1 Tracking Flights Approach

The following steps are taken to compare flight tracks and recommended routes:

1. Extract all fixes from the recommended route.
2. Extract, from FIX and Navaids datasets, the geographical coordinates of the fixes extracted at the previous step.
3. Store these coordinates in a list in the same order than the recommended route.

An algorithm and a set of metrics are then developed to assess the compliance of one flight to a recommended reroute.

4.4.2 Algorithm to compare flight trajectory and reroutes

To assess the compliance of flights to recommended reroutes, the Tracking Flights approach has been selected and is based on two sequences of coordinates: one for the recommended reroute and one for the flight trajectory. Two different approaches have been tested to assess the compliance of a flight to a reroute:

- A "polygon approach" based on the definition of a polygon around the recommended reroute and the presence of flight positions into this polygon.
- A "circle approach" based on the definition of circle areas around each element (waypoint, Navaid) of the reroute and the presence of one position of the flight within these circles.

4.4.3 Polygon approach

As mentioned before, this approach is based on the definition of a polygon, which can also be described as a corridor, around the reroute. The main parameter that needs to be defined for this approach is the width of the polygon. It is defined as 10 nautical miles (18.52 km) based on recommendations from FAA analysts.

In order to build the polygon around a reroute, the author created an algorithm that first builds a rectangle of ten nautical miles width and one nautical mile height around each element of the reroute. These rectangles are oriented to the next element of the reroute such that at the end the polygon is the contour linking all individual rectangles. The Figure 9 is an example of a recommended reroute between Newark Liberty International Airport and Fort Lauderdale-Hollywood International Airport issued November 6th, 2018.

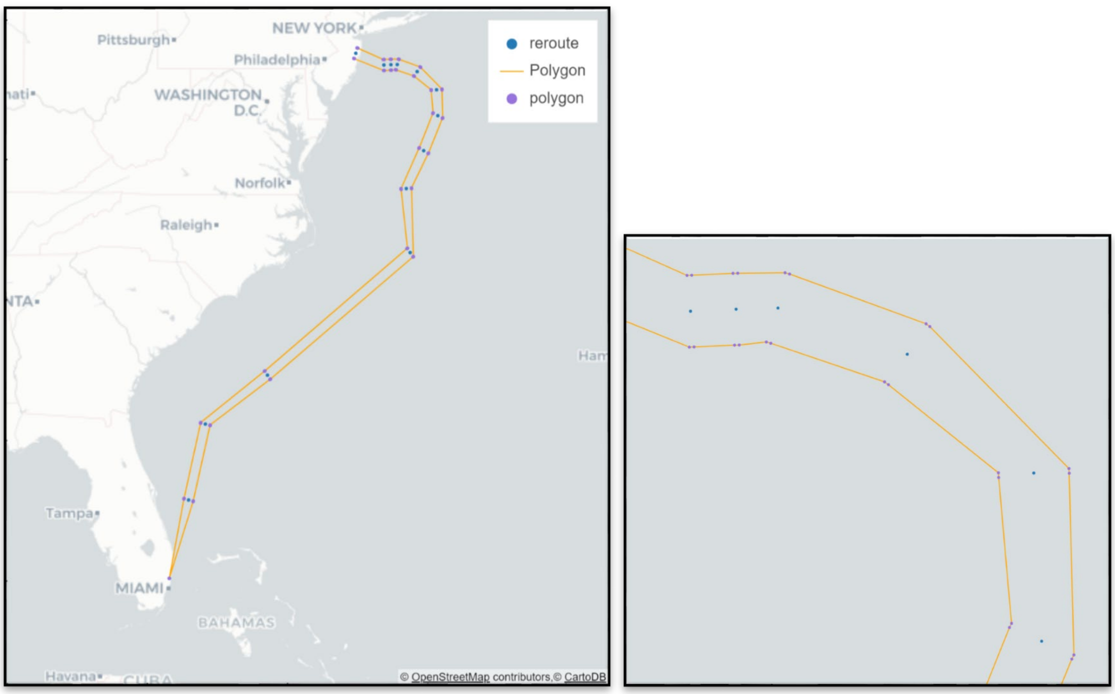

Figure 9: Polygon around recommended reroute EWR-FLL, November 6th, 2018

These visualizations are realized directly from Python in Bokeh [35] which creates html files with maps of the area of interest from the OpenStreetMap website. It is also possible to display the affected flights trajectories on these maps to visually assess their compliance to recommended reroutes, as illustrated in Figure 10.

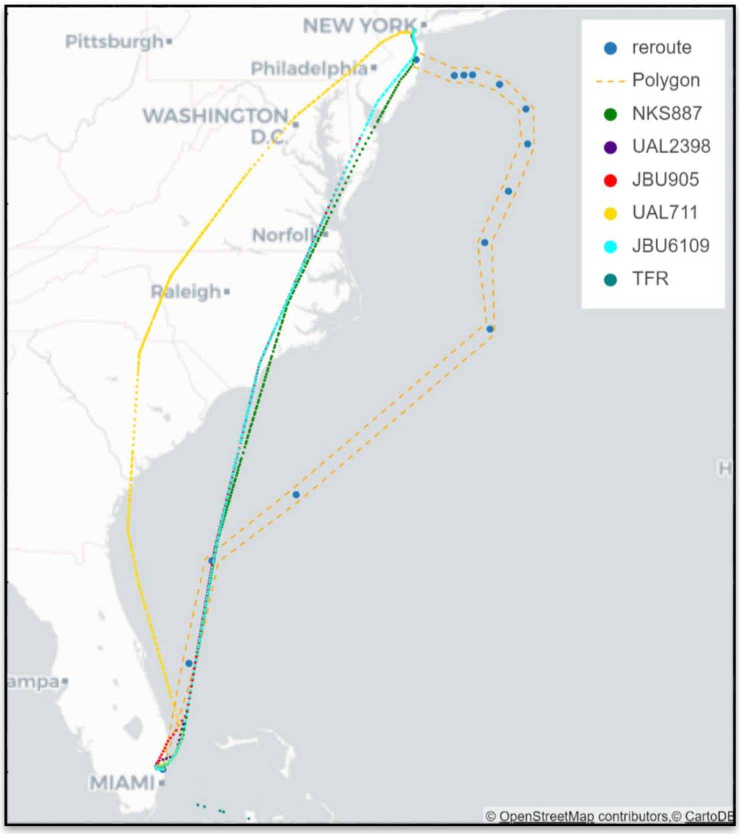

Figure 10: Recommended reroute EWR-FLL, Nov 6th 2018, and the affected flights trajectories

4.4.4 Circles approach

As mentioned before, the circle approach is based on the generation of one geographical circle around each element of the reroute. The compliance of a flight to the reroute is then measured according to the ratio of reroute elements that have at least one flight position within its circle and the total number of reroute elements. A reroute element that has at least one flight position within its circle is said to be "validated". The main parameter to fix for this approach is the radius of each circle around reroute elements. To be coherent with the polygon approach, this radius has been set to 5 nautical miles such that its diameter is the same length as the rectangle's width.

This approach is illustrated in Figure 9 for the same reroute advisory as previously, between Newark International Liberty Airport and Fort Lauderdale-Hollywood International Airport, on November 6th, 2018.

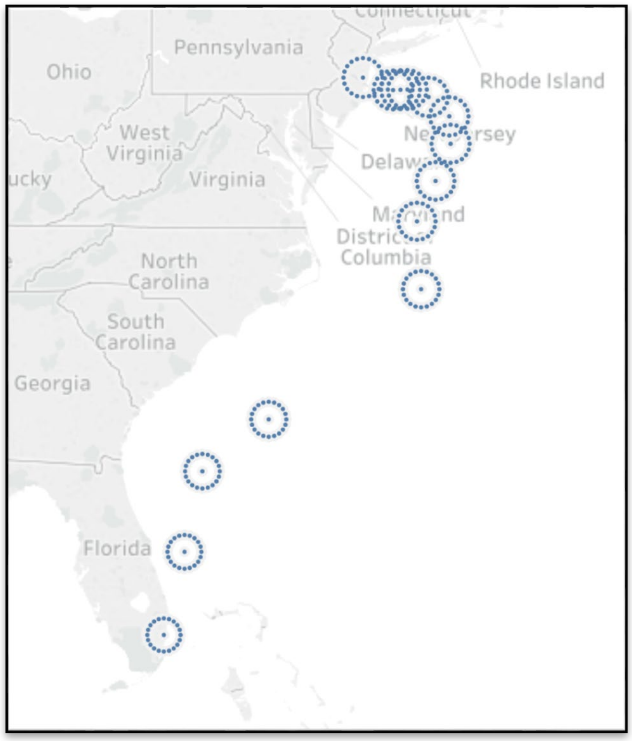

Figure 11: Circles around recommended reroute EWR-FLL, November 6th, 2018

4.4.5 Metrics

The two approaches defined above need to be completed with definitions of metrics that will measure the compliance of flights to reroutes. Four metrics have been developed and tested for the scope of this work.

4.4.5.1 Metric 1: polygon metric

The first metric to assess the compliance of flight to reroutes for the polygon approach is based on the following ratio:

$$metric\ 1 = \frac{Number\ of\ flights\ positions\ in\ the\ reroute\ polygon}{Total\ number\ of\ flights\ positions}$$

This metric is relevant if the flight's positions are uniformly distributed. However, most of the time, the frequency at which the plane issued its positions is higher during the ascent and descent phases of the flight. Thus, more metrics based on the polygon approach need to be introduced.

4.4.5.2 Metric 2: Flight distance metric

This metric also relies on the polygon approach detailed above but instead of dealing with the number of positions, it deals with the distance flown by the aircraft within the polygon and the total distance flown by the plane. It is defined as follows:

$$metric\ 2 = \frac{Distance\ flown\ by\ the\ plane\ within\ the\ polygon}{Total\ distance\ flown\ by\ the\ plane}$$

This metric is more reliable than the first one because it does not depend anymore on the frequency of positions issuance.

4.4.5.3 Metric 3: Reroute distance metric

This metric also relies on the polygon approach and more specifically on the length of the reroute, which is defined has the sum of distances between all waypoints of the reroute. It is defined as follows:

$$metric\ 3 = \frac{Distance\ flown\ by\ the\ plane\ within\ the\ polygon}{Length\ of\ the\ reroute}$$

This metric is close to the previous one but focuses more on the relation between the plane's trajectory and the reroute.

4.4.5.4 Metric 4: Circle metric

The metric to assess the compliance of flights to reroutes for the circle approach is based on the notion of "validated waypoint". A waypoint is said to be "validated" by a plane when it is possible to find at least one plane position within the circle of 10 miles diameter generated around the waypoint. Then the compliance of a flight to the reroute is measured by the following ratio:

$$metric\ 4 = \frac{Number\ of\ waypoints\ validated\ by\ the\ flight}{Total\ number\ of\ waypoints\ of\ the\ reroute}$$

Similarly to the previous metric, this metric is very efficient to measure flights compliance to reroutes if waypoints and navaids of the reroutes are uniformly distributed. However it appears, as shown in Figure 12, that most of the time, fixes and navaids are not distributed uniformly and that their frequency is usually higher at the beginning and end of the reroute, especially in the North-East area.

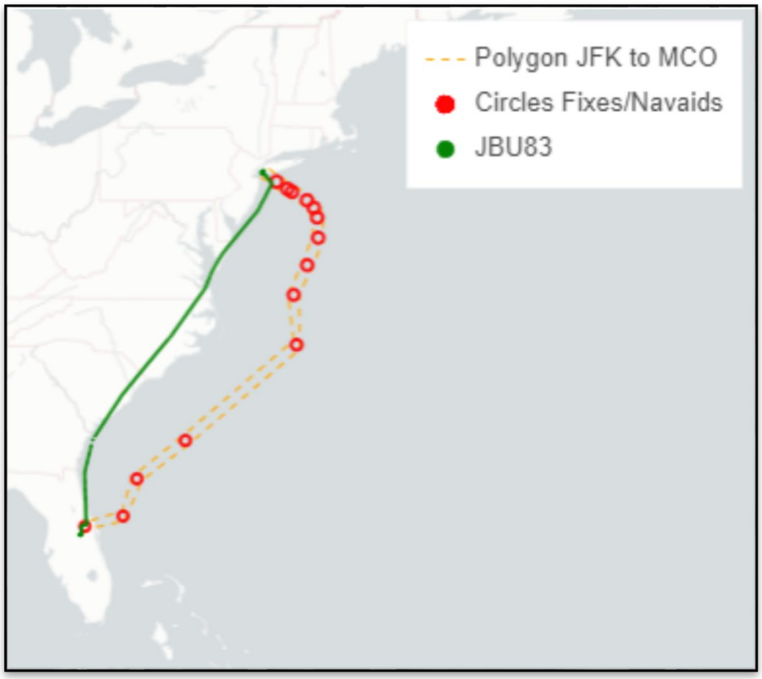

Figure 12: Example of reroute with unbalanced distribution of fixes and navaids

4.4.6 Evaluate the results

The aforementioned approach provided results about the compliance of pilots to these recommended reroutes and helped to identify relevant trends about that compliance.

The results about the compliance of flight operators to recommended flight reroutes help test *Hypothesis 1* and answer *Research Question 1*. *Hypothesis 2* and *Research Question 2* are addressed in the following section with the generation of a model to predict the issuance of volume-related reroute advisories.

4.5 Step #5: Generation and Validation of the Prediction Model

Reroute advisories due to volume are issued when Air Traffic Controllers notice or expect high traffic in a certain area of the national airspace. As discussed in Section 3.2, all reroute advisory types (Required, Recommended, FYI) issued between January and

April 2017 have been considered and extracted from Traffic Flow Management System (TFMS). Another key piece of information that needs to be collected in order to have a detailed situation of the air traffic in the national airspace is the hourly traffic count per facility. However, this dataset has not been provided by the FAA to the author because the time processing requested to extract the data was too long for this data to be included in the research. The data has therefore been generated by the author as logically as possible based on daily traffic counts that can be found online on the Air Traffic Activity Data System (ATADS) website owned by the FAA [36]. The distribution of flights per hour has been computed on July 27th, 2017 by the FAA [37] and is presented in Figure 13.

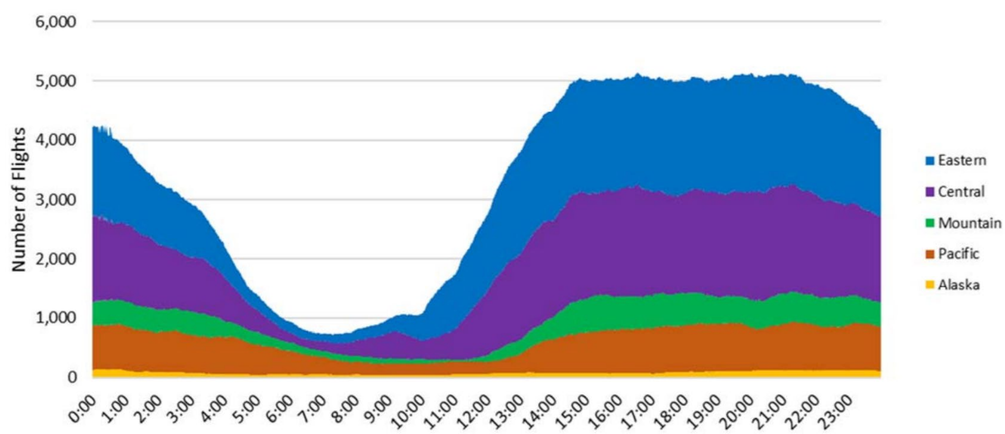

Figure 13: Number of flights in the US airspace per hour (GMT) and time zone on July 27th, 2017 [37]

The author used Figure 13 to compute the percentage of flights per hour which is plotted in the following diagram.

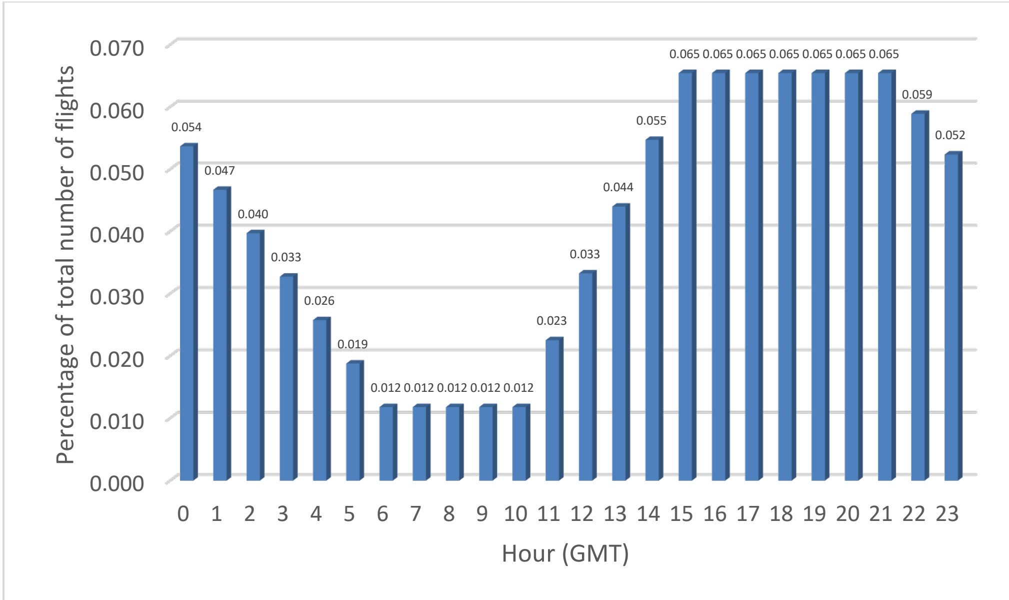

Figure 14: Distribution of the number of flights per hour (GMT) in the US airspace

The hourly traffic count per facility was then generated by multiplying the daily traffic count per facility with the hour percentage presented in the diagram above.

The prediction model presented in this work is therefore a proof of concept that can be easily converted to a realistic prediction model as soon as the actual hourly traffic count becomes accessible.

4.5.1 Process of Supervised Machine Learning

As stated in Section 2.2, supervised learning algorithms attempt to discover and model relationships between a target value and other values. This involves identifying and acquiring datasets, processing them into suitable formats, fusing the data, and training the data. The number of rows contained in the fused dataset is equal to:

$$4\ months * Nb\ of\ days\ per\ month * 24\ hours * 20\ ARTCC = 57600\ rows$$

However during this period, only 4511 rows in the dataset have been found to have one reroute issued. Therefore the predictor “Presence of a reroute” had the value ‘yes’ in 7.8% of cases and ‘no’ in 92.2% of cases. These ratios are too unbalanced to be processed correctly with Machine Learning algorithms. Therefore, the dataset has been under-sampled to a ratio between the two classes of 2-1. This ratio can be processed by machine learning algorithms. The data is not perfectly balanced but it preserves the initial nature of the dataset is preserved. Machine learning algorithms can then be applied on this reduced dataset. For the model to be properly evaluated, the data needs to be randomly divided into 3 subsets. The first one is the training dataset which represents half of the data and is used to generate the model. One fourth of the data is the validation subset which is used to iterate and refine the model. The remaining one fourth of the data is the test set used to evaluate the model.

The evaluation step comes after the validation of the model and is critical in the process as it informs how a learner will perform on future data.

4.5.2 Model generation

The model developed is based on a dataset issued by the fusion of TFMS dataset and traffic count data generated by the author. Fusing these two datasets will relate any volume-related reroute advisory with the actual traffic situation of the national airspace at any point in time. The generation of a predictive model able to predict the issuance of volume-related reroute advisory relies on the following steps:

1. Extract all volume-related reroute advisories issued in TFMS between January and April 2017
2. Collect hourly traffic counts per facility for the same time period
3. Build a data matrix gathering the extracted data from traffic counts and TFMS
4. Apply several supervised machine learning algorithms to predict the issuance of volume-related reroute advisories

4.5.3 Model Prediction

Seven Machine Learning techniques have been implemented in R and then tested on the fused dataset in order to determine which one performs best for this problem. The following two different predictions have actually been examined:

- Prediction of the issuance of volume-related reroute advisories without specifying the reroute type
- Prediction of the issuance of volume-related advisories with specifying the reroute type

For the first prediction, only two classes of prediction are possible:

- Yes, a volume-related reroute advisory has been issued
- No volume-related reroute advisory has been issued

For the second prediction, the model additionally tries to predict the reroute type. Thus, there is still the class ‘No reroute’ but also as many prediction classes as reroute types for the second prediction.

4.6 Step #6: Evaluation of the model

4.6.1 Confusion Matrix

Many different metrics exist to evaluate a Machine Learning model. Classification problems are typically evaluated using results of a confusion matrix, a performance measurement tool for supervised learning classification problems where outputs are two or more classes. More precisely, a confusion matrix is a table gathering four different combinations of predicted and actual values and which results are used in most of performance metrics such that Recall, Precision, Specificity, Accuracy, etc. [38]

Table 3: Confusion Matrix

		Actual Values	
		Positive	Negative
Predicted values	Positive	True Positive (TP)	False Positive (FP)
	Negative	False Negative (FN)	True Negative (TN)

- True Positives (TP) are the cases when the predicted positive values are also actual positive.
- False Positives (FP) are the cases when the predicted values are positive while they are actually negative.
- True Negatives (TN) are the cases when the model predicted negative values and the actual values are also negative.
- False Negatives (FN) are the cases when the model predicted positive values while the actual values are negative.

4.6.2 Performance metrics

Various algorithms are tested in order to find the most accurate one to predict the issuance of a reroute advisory due to volume constraint but also the most accurate one to predict the type of the potentially issued reroute advisory. Performance metrics detailed below are used to compare those algorithms and isolate the best performing one.

Accuracy [39]

The accuracy in classification problems is the total of correct predictions (True Positives + True Negatives) over the total of all predictions.

$$Accuracy = \frac{TP + TN}{TP + FP + FN + TN}$$

Accuracy is very appropriate to use when the two data classes to predict are numerically nearly balanced. However, and accordingly to what has been stated in Section 4.5.1, the dataset used is not balanced, which implies that the Accuracy is not one of the main metric used to compare the various algorithms.

Precision [39]

Precision refers to the ratio of correct positive predictions over the total of all positive predictions:

$$Precision = \frac{TP}{FP + TP}$$

Sensitivity [39]

Sensitivity measures the number of correct positive predictions over the total number of actual positive values that should have been predicted if the model was perfect.

$$Sensitivity = \frac{TP}{TP + FN}$$

Specificity [39]

Specificity measures the proportion of correct negative predictions over the actual number of negative values. Specificity is the exact opposite of Sensitivity.

$$Specificity = \frac{TN}{TN + FP}$$

Matthew's Coefficient [40]

Matthew's Coefficient's maximum value is 1 and corresponds to perfect predictions for the test dataset. Matthew's Coefficient's minimum value is -1 and corresponds to total contradiction. It is specified as:

$$MCC = \frac{TP * TN - FP * FN}{\sqrt{(TP + TN) * (TP + FP) * (FP + TN) * (TN + FN)}}$$

The MCC is a relevant metric to compare the algorithms performances for the first prediction because this prediction has only two classes: Yes Reroute, No Reroute.

Kappa Statistic

The Kappa statistic is very often used to measure the performance algorithms on multi-class and imbalanced datasets. It is defined with the following expression:

$$\kappa = \frac{p_0 - p_e}{1 - p_e}$$

Where p_0 is the observed value and p_e is the expected value. According to J.R. Landis and G.G. Koch [41], values of Kappa statistic is an indicator of the performance of the classifier. Following is their interpretation of some Kappa statistics' values:

Table 4: Interpretation of Kappa statistic from Landis and Koch, 1977

Kappa Statistic	Strength of agreement
<0.00	Poor
0.00 – 0.20	Slight
0.21 – 0.40	Fair
0.41 – 0.60	Moderate
0.61 – 0.80	Substantial
0.81 – 1.00	Almost Perfect

The Kappa statistic is a very relevant comparison metric to be used for both prediction since the dataset issued from the fusion of reroute advisories and traffic counts is not perfectly balanced. Indeed, unlike the accuracy metric, the Kappa statistic is well suited to assess the performance of models based on unbalanced data. Chapters 5 and 6 present the results obtained from implementing the approach discussed in Chapter 4.

CHAPTER 5. QUANTITATIVE ASSESSMENT OF THE RELEVANCY OF RECOMMENDED FLIGHT REROUTES

5.1 Data acquisition and processing

5.1.1 Traffic Flow Management System (TFMS)

The Federal Aviation Administration (FAA) provided the author with one year (2017) of TFMS Flow data containing all advisories that were issued as part of Traffic Management Initiatives. The 2017 year represents two terabytes of data that needed to be processed in order to extract all recommended flight reroutes. This work has been achieved over a period of four months (January to April 2017) following the methodology described in the previous chapter.

Each recommended reroute message in TFMS provides the reroute as a sequence of waypoints, Navaids and airways numbers. In order to compare flights trajectories and reroutes, they have been converted from a sequence of waypoints and Navaids to a sequence of GPS positions. These positions are found in datasets published and updated every 28 days by the FAA [31].

5.1.2 System Wide Information Management Flight Data Publication Service (SFDPS)

The Federal Aviation Administration (FAA) provided the author with ten days of SFPDS data containing information about all flights flying over the US territory. The SFDPS dataset is significantly larger than TFMS: one day of SFDPS data is always around 100 gigabytes while one day of TFMS data is around 500 megabytes. These ten days of data have been processed by the author according to the methodology described in the previous chapter. TH messages issued by flights during these ten days have been extracted and stored as CSV files. The rest of the data needed has been processed directly by the FAA and sent as CSV files to the author.

5.2 TFMS and SFDPS fusion

The purpose of fusing SFDPS and TFMS datasets is to extract from SFDPS the flights affected by recommended reroutes gathered in TFMS. This fusing step has been achieved using Python and more specifically the following modules:

- Csv module [42]: used to read and write csv files.
- Pandas module: used to process large csv files.
- Datetime module: used to read dates and times, specifically to make sure a flight is in the time window of the recommended reroute.
- Math module: used for all mathematical operations needed such as conversion between geographical coordinates formats (GPS, DMS).
- Numpy module: used for manipulations of matrices and lists.
- Matplotlib.path module: used for generation of polygons around trajectories.
- Bokeh module: used for visualization of trajectories and reroutes on geographical maps.

According to the methodology described in the previous chapter, extracted flights are those that have their origin and destination airport captured in the scope of the reroute and that are flying during the effective period of the reroute. For each affected flight a list is created in order to store its trajectory. Because TH messages are issued every twelve seconds, only one message out of five is considered such that the trajectory of each affected flight is made of positions updated every minute.

5.3 Results

Results presented below are based on the analysis of four months of data corresponding to January 1st 2017 to April 30th 2017. During this period 4,974 recommended reroutes affecting flights have been recorded and 22,016 flights affected by these reroutes have been analyzed.

Measuring the performance of each metric with real examples is necessary to determine the one(s) that is (are) the most suitable and should be used to assess the compliance of flights to recommended reroutes.

5.3.1 Metrics performance

Metric 1

$$metric\ 1 = \frac{Number\ of\ flights\ positions\ in\ the\ reroute\ polygon}{Total\ number\ of\ flights\ positions}$$

As mentioned earlier, the metric 1's performance is very related to the frequency of flight positions issuance. If the distribution is not uniform, then the ratio is biased and does not represent the compliance of the flight to the reroute.

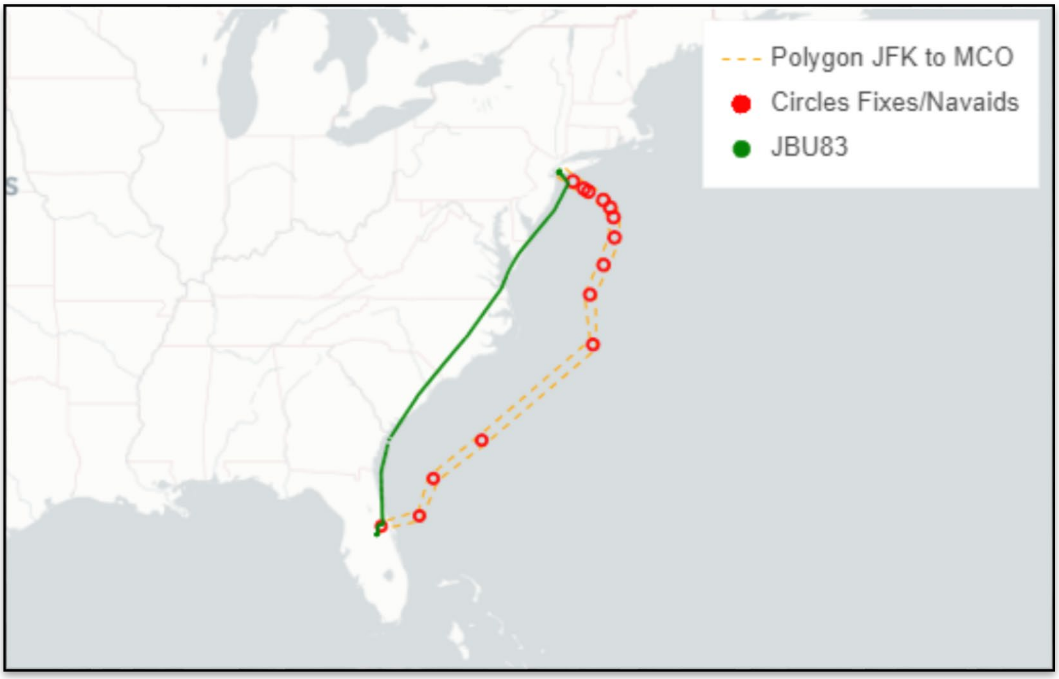

Figure 15: Example of flight with corrupted metric 1

For example, flight JBU83 in Figure 15 is said to follow the reroute with a compliance of 43% based on metric 1 but 0.07% based on metric 2. It is clear that metric 2 is much more accurate than metric 1 for this example. Metric 1 is biased because the frequency of positions issuance during the ascent and descent phases is much higher than during the rest of the flight. This a typical example of flight where the distribution of position is not uniform.

Metric 2

$$metric\ 2 = \frac{Distance\ flown\ by\ the\ plane\ within\ the\ polygon}{Total\ distance\ flown\ by\ the\ plane}$$

Metric 2, because it relies on the distance flown by the plane and not the frequency of positions issuance demonstrated better performance than metric 1. When the flights data is cleaned and the polygon well defined, then this metric appears to be very reliable and no situation has been found where it is biased.

Metric 3

$$metric\ 3 = \frac{Distance\ flown\ by\ the\ plane\ within\ the\ polygon}{Length\ of\ the\ reroute}$$

Metric 3 is accurate for most of the cases because it is not dependent on the position's distribution. However, one case has been found where it is biased. Such instance is due to the format of the reroute defined as the "second format" in the Chapter 4. This format is based on the definition of route origin segments and route destination segments. As explained in Chapter 4, it happens sometimes that the origin and destination segments overlap and thus might create back and forth portion in the reroute. The following Figure 16 shows an example.

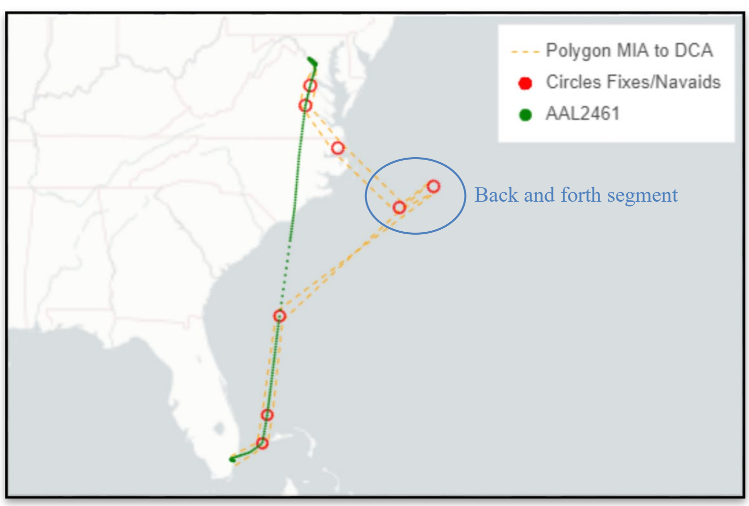

Figure 16: Example of flight and reroute with corrupted Metric 3

For this example, the following values of the metrics have been computed:

$$metric\ 2 = 0.501$$

$$metric\ 3 = 0.380$$

For this example, the recommended reroute between Miami International Airport and Ronald Reagan National Airport has a problem that affects the performance of metric 3. Indeed, the flights are asked to take a portion of the reroute (between waypoints 4 and 5) back and forth. Therefore, the length of the reroute is longer than what it should be and thus, the ratio of metric 3 is lower. Without the back and forth portion, the ratio of metric 3 would be 0.45 which is closer from the value of metric 2.

This problem in the definition of recommended reroute is recurrent and is due to the way Air Traffic Controllers define recommended reroutes. It could be solved by a deeper analysis of the reroute with an algorithm that would detect and delete these back and forth portions. This would ultimately improve the performance of metric 3.

Metric 4

$$metric\ 4 = \frac{Number\ of\ waypoints\ validated\ by\ the\ flight}{Total\ number\ of\ waypoints\ of\ the\ reroute}$$

Similarly to metric 1, the metric 4 is very sensitive to the distribution of the reroute waypoints. Figure 17 shows an example where metric 4 is biased because of the large numbers of waypoints in the New York Area that are not validated and the few numbers of validated waypoints where the plane follows the reroute.

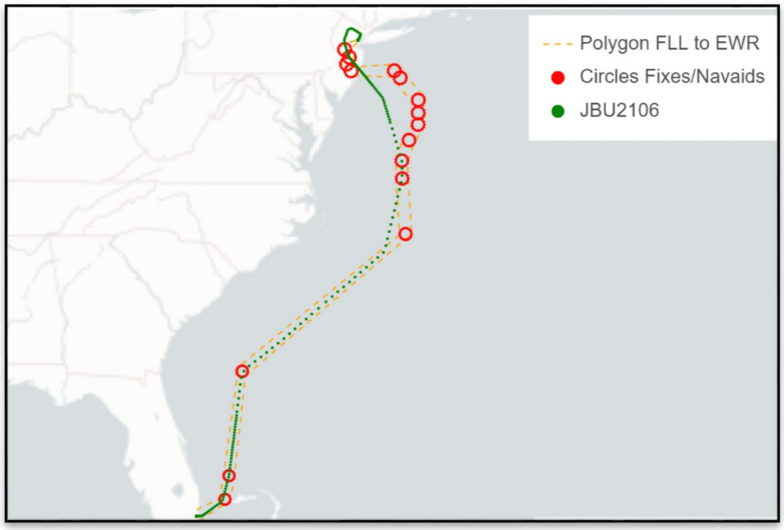

Figure 17: Example of flight and reroute with corrupted metric 4

For this example, the following values of the metrics have been computed:

$$metric\ 2 = 0.70$$

$$metric\ 3 = 0.67$$

$$metric\ 4 = 0.44$$

This pattern is very recurrent in the North East area of the National airspace because of the very important traffic around New York, Boston, Philadelphia and Washington DC. Indeed, in these very dense areas, Air Traffic Controllers provide very precise path for reroutes which means a lot of waypoints.

According to those considerations on metrics performances, it appears that situations have been found where metrics 1, 3 and 4 are biased and therefore can lead to errors in the assessment of flights compliance to recommended reroutes. Metric 2, computing the ratio of the distance flown by the plane within the polygon and the total distance flown by the plane, is the most reliable and accurate one since no instance has been found where it is biased. Thus, metric 2 is the one used to assess the compliance of flights to recommended reroutes. After those considerations on metrics' performance, it is possible to analyze more precisely the compliance of flights to recommended reroutes.

5.3.2 Compliance results

All 22,016 flights affected by a recommended reroute between January and April 2017 have been processed and their compliance to the recommended reroute have been measured based on metric 2 detailed above. The distribution of flight according to metric 2 is detailed in Figure 18.

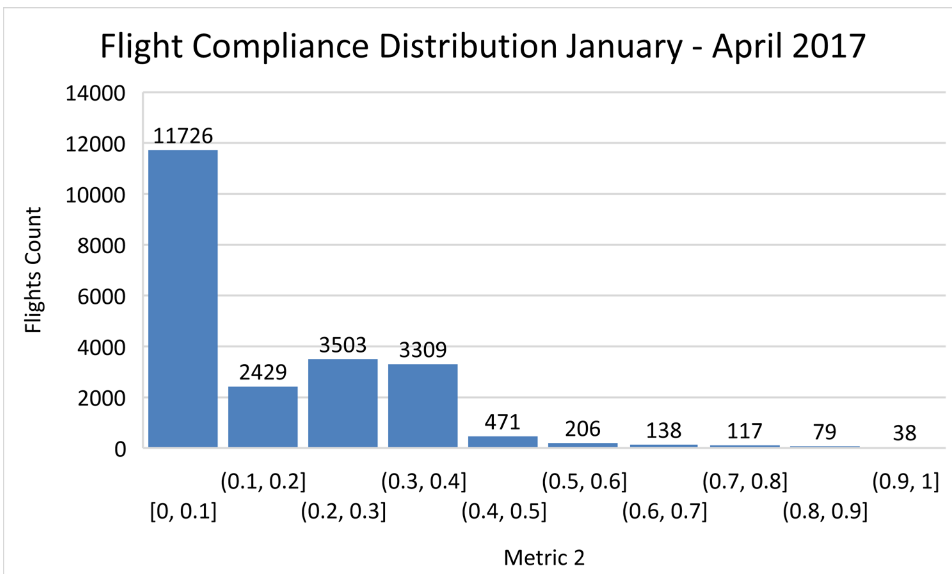

Figure 18: Flight compliance to recommended reroutes based on metric 2 for the period January – April 2017

Based on this distribution of flights, it is interesting to define a threshold above which flights are considered to have followed the reroute. To do so, the best way is probably to visualize flights for certain values of metric 2 and visually determine if flights complied with the reroute. Following are some examples of flights with different values of metric 2.

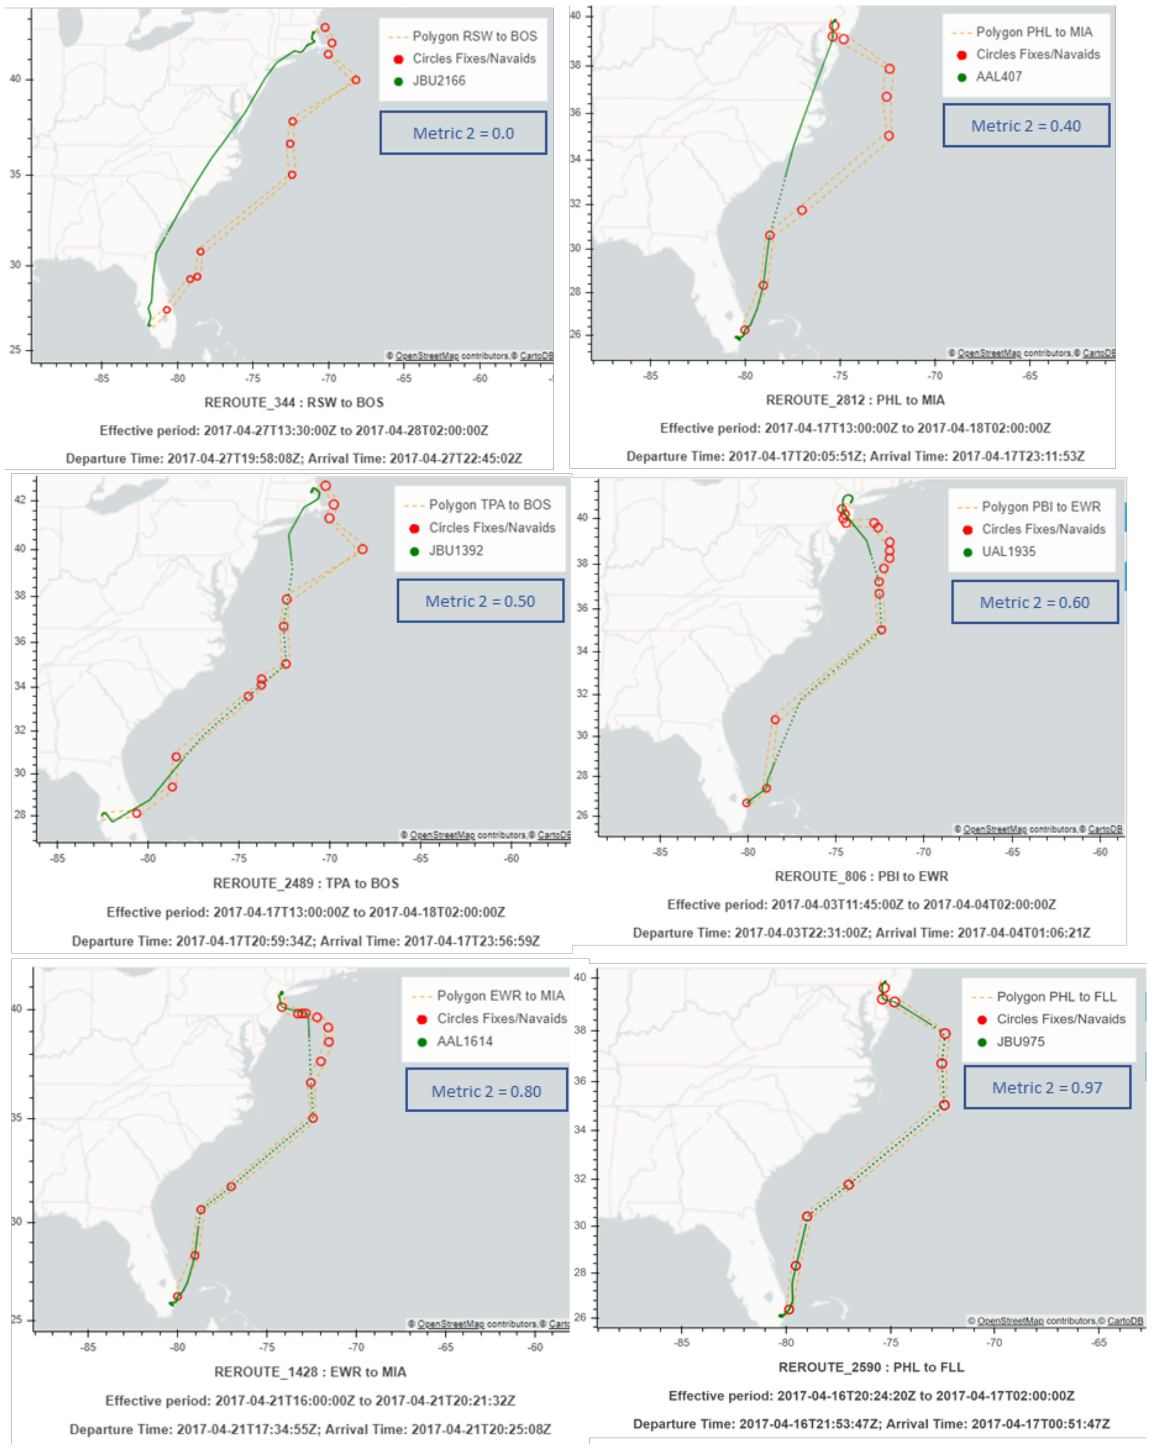

Figure 19: Visualization of flights and reroutes for different values of metric 2

Based on these examples and on the 22,016 flights analyzed, it appears that the threshold of compliance is located between the values of 0.5 and 0.6 for metric 2. Looking more deeply into values of metric 2 between 0.5 and 0.6, it appears that 0.55 can be considered as a good assumption for the compliance threshold. Following are two examples of flights with metric 2 values of 0.55 and 0.53 which represent well the general behavior of flights with such metric 2 values.

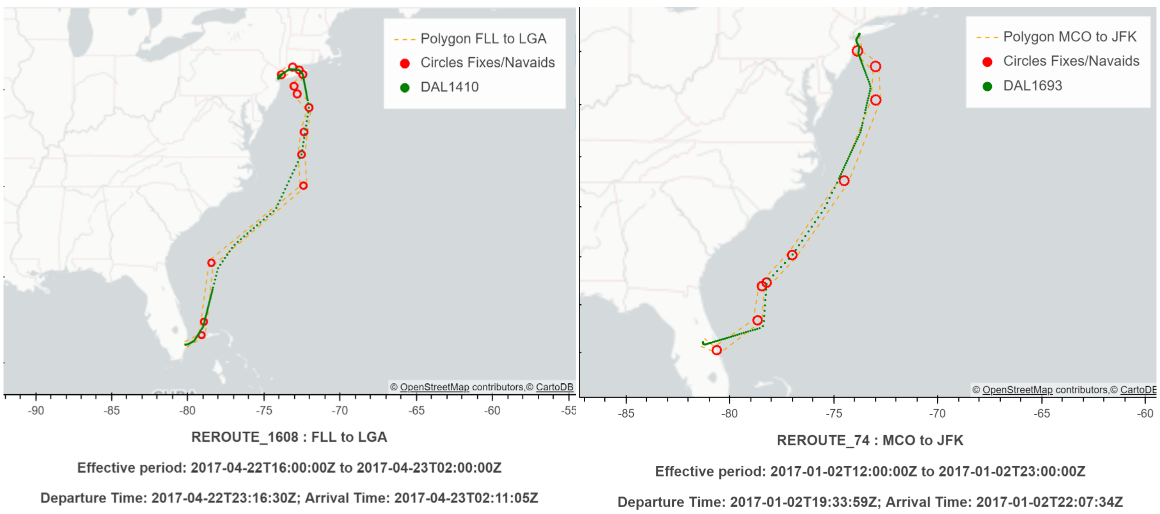

Figure 20: Examples of flights with metric 2 ratio of 0.55 (left) and 0.53 (right)

Table 5 shows the amount of flights that can be assumed to comply with reroutes and those who do not, with a threshold value of 0.55.

Table 5: Compliance of flights with a compliance threshold of 0.55

Metric 2	Count of flights	Percentage
<0.55	21559	97.9%
0.55-1	457	2.1%

Figure 21 shows the evolution of the flight compliance to recommended reroute with a compliance threshold varying from 0.3 to 0.9.

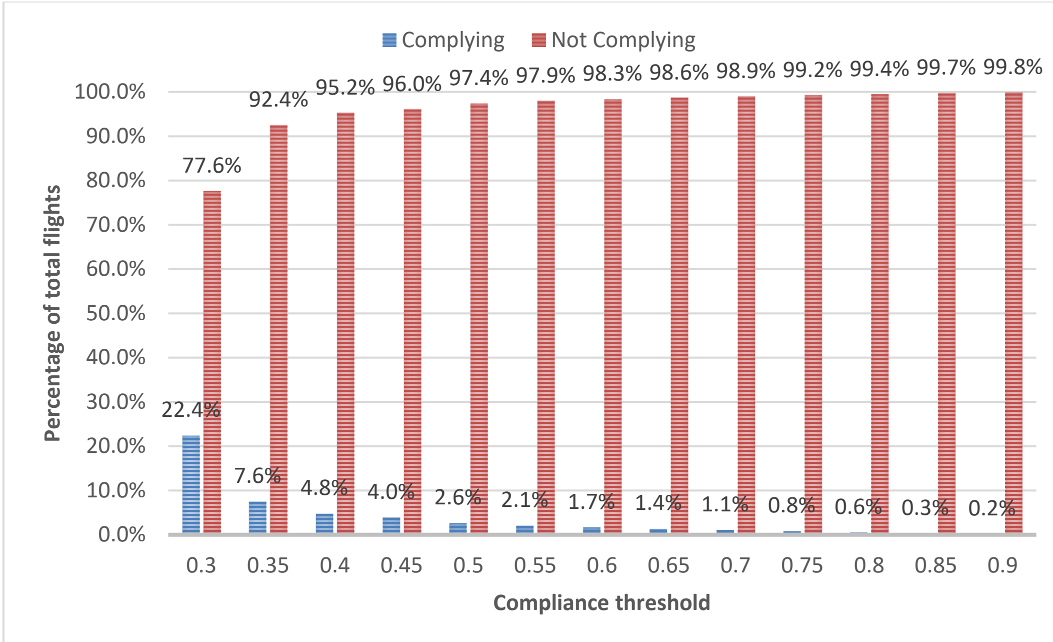

Figure 21: Evolution of flights compliance to reroute with the compliance threshold

As expected, and in accordance with Figure 19, the percentage of flights complying with recommended reroutes drops as the compliance threshold increases between 0.3 and 0.9. It is however noticeable that this compliance percentage drops twice as fast between thresholds of 0.3 and 0.35 (22.4% to 7.6%) than it does between thresholds of 0.35 and 0.9 (7.6% to 0.2%). It means that there are twice more flights with metric 2 in the range [0.30 ; 0.35] than with metric 2 in the range [0.35 ; 0.90].

5.3.3 Reroute Distribution per region

As mentioned in the first chapter, there are twenty-one facilities managed by Air Route Traffic Control Centers (ATRCC) in the National Airspace, distributed as shown in Figure 22.

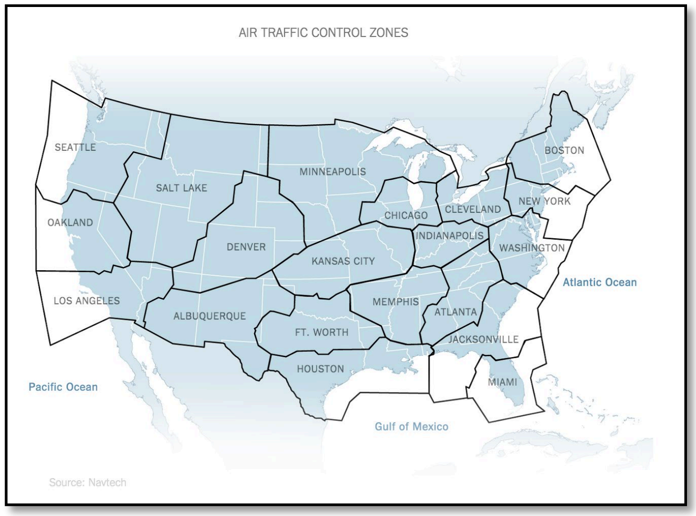

Figure 22: Air Traffic Control Sectors of the NAS [3]

Each recommended reroute is defined by the affected origin and destination airport. It is therefore possible to rank the NAS facilities according to the count of recommended reroutes they are affected by. This ranking is presented on the following heatmap (Figure 23).

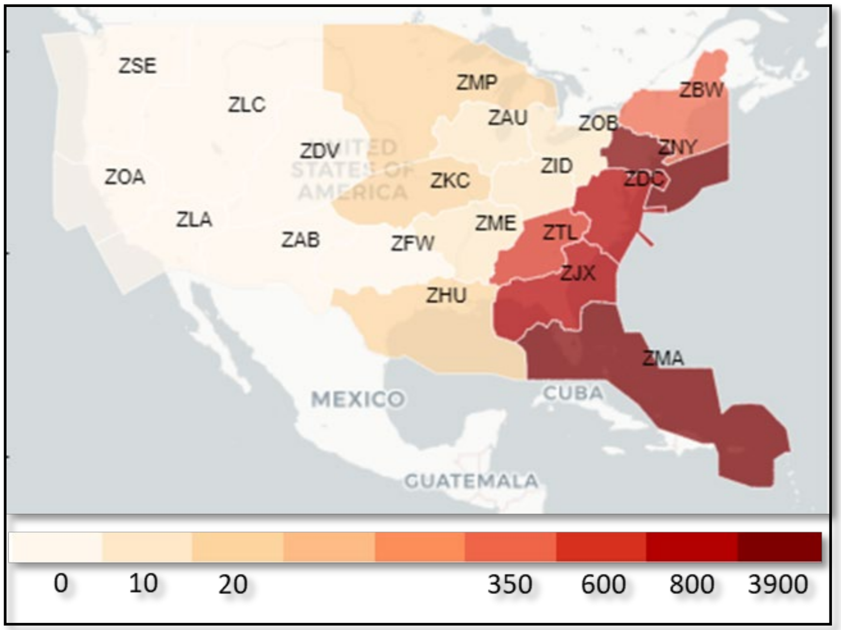

Figure 23: Heatmap of the count of reroutes issued per facility between January and April 2017

Figure 23 shows that a large majority of the recommended reroutes issued in the period between January and April 2017 are affecting the East coast of the United States. The most frequent connections between facilities are displayed in Table 6.

Table 6: Connections affected the most frequently by recommended reroutes

Routes	Number of reroutes issued
ZMA – ZNY	2770
ZMA – ZDC	513
ZJX – ZNY	468
ZMA - ZTL	293
ZMA - ZBW	260
ZNY – ZTL	220

Routes between Florida (ZMA, ZJX) and North-East (ZNY, ZDC, ZBW) represent 75% of all recommended reroutes issued between January and April 2017. No recommended reroutes has been issued on the West coast during this period. The difference in the number of recommended reroutes between the East and West coasts is mainly due to the very high distribution of important airports (New York area, Boston, Washington DC, Atlanta) on the East coast as well as the high traffic because of the connections with Europe.

The average value of metric 2 ratio on all flights analyzed without any consideration for the location of the origin and destination airports is 0.15. A study has been conducted on the flights with at least the origin or destination airport outside the East Coast region. The East Coast region has been delimited on the West side by a North-South line on Atlanta, Georgia and considering that Atlanta is part of the East Coast region. This subset consists of 83 flights with an average value of 0.51 for metric 2. This number shows first that flights outside of the East Coast region comply more with recommended reroutes than the ones on the East Coast. Furthermore, it can also be interpreted as the fact that recommended reroutes outside the East Coast region are much more relevant and efficient than those affecting connections within the East Coast region.

5.3.4 Flight compliance analysis according to distance flown

All 22,016 flights analyzed can be filtered according to the distance flown, hence allowing to capture the distribution of flights impacted by recommended reroutes for various ranges of distance flown.

Table 7 displays the number of flights for various distance ranges as well as the average value for metric 2 for each distance range.

Table 7: Number of flights impacted by Recommended Reroutes with average and median metric 2 per distance flown

Distance (km)	Distance (nm)	Count of flights impacted by RMD	Avg of Metric 2	Median Metric 2
0-500	0 - 270	35	0.17	0.15
500-1000	270 - 540	2473	0.06	0.03
1000-1500	540 - 810	2996	0.09	0.07
1500-2000	810 - 1080	14165	0.18	0.15
2000-2500	1080 - 1340	2259	0.18	0.07
2500-3000	1340 - 1620	59	0.16	0.02
3000-3500	1620 - 1890	18	0.03	0.02

According to the Table 7, it appears that most of flights (64%) analyzed are flying a distance between 1500km and 2000km. The table also shows that the average value of metric 2 is maximum (0.18) for distances varying between 1500 to 2000 km and 2000 to 2500km. However the median value of metric for distances between 2000 to 2500 km is low (0.07) compare to distances between 1500 to 2000km. Finally we can assume from this table that Air traffic Controllers define the recommended reroutes advisories mainly for mid-distance flights (1500 to 2000km) and these advisories are in general not followed for longer or shorter distances.

5.3.5 *Flight compliance analysis according to the type of airline*

The SFDPS dataset records flight information for all domestic flights operated by civil aircraft travelling in the National Airspace. These civil aircraft can be divided into three categories: legacy airlines, regional airlines and charter airlines (leasing private jets). Table 8 provides all American airlines and their ICAO identifier which is always the three first letters of the flight Id.

Table 8: Table of airlines of the United States

Legacy		Regional		Charter	
Airline	ICAO	Airline	ICAO	Airline	ICAO
Alaska Airlines	ASA	Air Wisconsin	AWI	JetSuite	RSP
Allegiant Air	AAY	Cape Air	KAP	NetJets	EJA
American Airlines	AAL	CommutAir	UCA	XOJET	XOJ
Delta Air Lines	DAL	Compass Airlines	CPZ	FlexJet	LXJ
Frontier Airlines	FFT	Contour Aviation	VTE	Delta Private Jets	DPJ
Hawaiian Airlines	HAL	Elite Airways	MNU	Hop-a-Jet	HPJ
JetBlue Airways	JBU	Endeavor Air	EDV	…	…
Southwest Airlines	SWA	Envoy Air	ENY		
Spirit Airlines	NKS	ExpressJet	ASQ		
Sun Country Airlines	SCX	GoJet Airlines	GJS		
United Airlines	UAL	Horizon Air	QXE		
		Mesa Airlines	ASH		
		PenAir	PEN		
		Piedmont Airlines	PDT		
		PSA Airlines	JIA		
		Republic Airline	RPA		
		Silver Airways	SIL		
		SkyWest Airlines	SKW		

The results obtained about the compliance of flights to recommended reroutes can be filtered according to the type of airline as shown in Table 9.

Table 9: Flights compliance according to the airline type

Airline Type	Count	Percentage	Median Distance flown (km)	Mean Metric 2	Median Metric 2
Legacy airlines	17734	78.70%	1710	0.156	0.09
Regional airlines	937	4.30%	1282	0.066	0.03
Charter – Private	3345	15.20%	1831	0.155	0.11

Table 9 shows that the vast majority of analyzed flights are operated by legacy airlines such as Delta Air Lines or American Airlines. It is interesting to notice that legacy airline flights and charter flights behave similarly towards recommended flight reroutes with a compliance ratio around 0.155. It is also remarkable that very few flights impacted by recommended reroutes are operated by regional airlines (4.3%) but those few have a very low average compliance ratio (0.066). Median values of metric 2 are lower than average but the same pattern is observable with a significant lower median value for regional airline flights. The median distance flown per type of airline shows that regional airlines flights are flying in general lower distances than legacy and charter airlines. These last remarks corroborate the assumption made in the previous section (*5.5.4*) stipulating that recommended reroutes seem to be designed for mid-distance flights (1500 to 2000km) but not for shorter one such as those operated by regional airlines.

5.4 Chapter conclusion

Four months of air traffic data have been processed and analyzed with Data Analytics methods such as Data Fusion in order to assess the efficiency and relevance of recommended flight reroutes. Once the two main sources of data were parsed and fused, the author developed a Python algorithm to compare flights trajectories and reroute paths and determine how much the flight did comply with the recommended reroute. Two approaches and four metrics were tested for this work and one has been isolated as the best performing one. This particular metric relies on the distance flown by the plane within a polygon built around the reroute path. According to this metric, only 2.1% of the 22,016 flights analyzed complied with the recommended reroute advisory they were affected by. This statistic shows that most of the time, pilots decide not to comply with recommended reroutes. It probably means that pilots consider these reroutes to deviate too much from the preferred route which impacts significantly other parameters, such as fuel consumption or the time delay.

Further analyses about different characteristics of flights such as the geographical area impacted by reroutes were conducted and showed for example that flights above the Midwest area are much more likely to comply with reroutes than those flying above the East Coast. No recommended reroutes affecting the West Coast of the National Airspace have been issued during the period analyzed by the author. It is coherent with general information about air traffic over the United States that indicate that the East Coast is much more impacted by volume constraints than the West Coast.

The work presented in this chapter has therefore provided a methodology to assess the relevancy of recommended reroutes and compute the compliance of flights to those reroutes over a certain period of time. This work thus provides an answer to the **Research Question #1** and validates the **Hypothesis #1** stipulating that Data Fusion can be used to analyze recommended reroute messages and flight data to identify trends and assess the relevancy of recommended flight reroutes. The next Chapter details the work conducted to answer the Research Question 2 about the prediction of volume-related reroute advisories.

CHAPTER 6. PREDICTIONS ON VOLUME - RELATED REROUTE ADVISORIES

The objective of the model detailed in this chapter is to predict the issuance and the type of volume-related reroute advisories. As stated Section 4.5 of the Methodology Chapter, this model is a proof of concept since it has not been possible to access hourly traffic count data during a long enough period of time. Seven different Machine Learning algorithms known for their classification abilities have been benchmarked to identify the best performing one on this problem: Decision Tree, Nearest Neighbor, Naïve-Bayes, Support Vector Machines, Bagging Ensembles, Boosting Ensembles, and Random Forest. The following two different predictions have actually been examined:

1. Prediction of the issuance of volume-related reroute advisories without specifying the reroute type
2. Prediction of the issuance of volume-related advisories with specifying the reroute type

The algorithms have been implemented in R and to make sure that they were compared accurately, the data was always randomly separated into three subsets used for each step of the implementation: training set, validation set, and testing set. The predictors for this model are the traffic count, the facility, the month, the day, and the hour. For the first prediction, models can be evaluated using Kappa statistic and Matthew's Coefficient because there are only two classes of prediction: No Reroute and Yes Reroute. However, if the type of the reroute is also predicted then Matthew's coefficient is not usable anymore because there is still the class No Reroute but also as many classes as reroute types. Thus, for the second prediction, the main performance metric to be analyzed is the Kappa statistic since the dataset is unbalanced (see Section 4.5.1). The analysis of the dataset issued from the fusion of volume-related reroute advisories and traffic count data showed that only Recommended and Required reroute advisories have been issued because of volume constraints between January and April 2017. Thus, the second prediction has three classes: No Reroute, RQD Reroute and RMD reroute.

6.1 Decision Trees

The following steps were taken to implement Decision Trees algorithm in R:

1. Load the input csv data using the R function "read.csv"
2. Randomly divide the data into training, validation and testing sets
3. Use the "C50" function to train the model on the training set
4. Test the performance of the model on the validation set with "predict" function
5. Use adaptive boosting to improve the performance of the dataset. This consists in building multiple decision trees and adding a "trials" parameter such that the optimal number of trials produces the lowest number of incorrect predictions.
6. Run the model on the testing set and build confusion matrices for validation and testing set with the "confusionMatrix" function.

6.1.1 *Prediction of the issuance of volume-related reroute without specifying its type*

Table 10 provides the confusion matrix for the first prediction on the testing set. The model had an accuracy of 0.9279, a Kappa statistic of 0.8391, a MCC of 0.839, a specificity of 0.9425 and a sensitivity of 0.8991.

Table 10: Confusion Matrix of Decision Tree predictions on the issuance of volume-related reroute advisories

	Actual No Reroute	**Actual Yes Reroute**
Predicted No Reroute	2114	115
Predicted Yes Reroute	129	1025

Table 10 shows that the model predicted correctly 2114 No Reroute instances and 1025 Yes Reroute instances. However, it also predicted incorrectly 115 Yes Reroute and 129 No Reroute. Both sensitivity and specificity are high for a classification problem but the fact that the sensitivity is higher than specificity demonstrates that the model is better at predicting the No Reroute class than the Yes Reroute.

The implementation of Decision trees algorithm with the “C50” function in R allows to extract the ranking of predictors according to their importance in the model. Figure 24 plots this importance per predictor. As expected, the traffic count and ARTCC sectors are the two most important predictors. The Month, Day, Hour predictors may gain importance if the model was trained on two years’ worth of data.

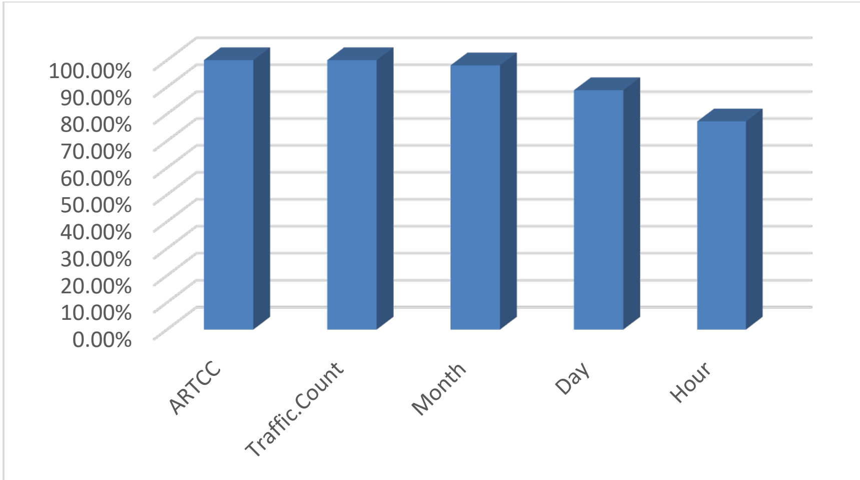

Figure 24: Importance of predictors in the Decision Tree model

The overall performance of Decision Tree technique for the first prediction is measured by both Kappa statistic (0.8391) and MCC (0.839). Based on these values and on the interpretation of Landis and Koch [41], the Decision Tree model can be considered as an almost perfect classifier for this prediction.

6.1.2 Prediction of the issuance of volume-related reroute with specifying its type

Table 11 provides the confusion matrix for the prediction of the issuance of volume-related reroute advisories with specifying the reroute type. The model had an accuracy of 0.9231 and a Kappa statistic of 0.8356.

Table 11: Confusion Matrix of Decision Tree prediction on the issuance and the type of volume-related reroute advisories

	Actual No Reroute	Actual RMD Reroute	Actual RQD Reroute
Predicted No Reroute	2141	14	106
Predicted RMD Reroute	3	100	11
Predicted RQD Reroute	99	27	882

Table 11 shows that the model correctly predicted 2141 No Reroute situations but incorrectly predicted 14 RMD Reroute advisories, and 106 RQD Reroute advisories. It also predicted correctly 100 RMD Reroute advisories, and incorrectly 11 RQD Reroute advisories and 3 No Reroute situations. Finally, it predicted correctly 882 RQD Reroute advisories but incorrectly 27 RMD Reroute advisories, and 99 No Reroute situations.

Table 12 summarizes the Sensitivity and Specificity metrics computed for all classes of the Decision Tree model. The high sensitivity and specificity values of the No Reroute and RQD reroute classes show that the model predicted well the actual issuances of required reroute advisories without making many mistakes. However, the significantly lower sensitivity of RMD Reroute class show that the model is not as good at predicting the issuance of recommended reroute advisories.

Table 12: Sensitivity and Specificity metrics for each class of the prediction model with Decision Tree technique

	Class: No Reroute	**Class: RMD Reroute**	**Class: RQD Reroute**
Sensitivity	0.9545	0.70922	0.8829
Specificity	0.8947	0.99568	0.9471

The overall performance of the Decision Tree algorithm for the second prediction is measured with the Kappa statistic of 0.8346. Based on this value and on the interpretation of Landis and Koch [41], and as for the first prediction, the Decision Tree model can be considered as an almost perfect classifier for the second prediction.

6.2 k-Nearest Neighbor (k-NN)

This technique is one of the simplest classification technique because it does not make any assumptions on the data and obtain high accuracy in general. According to the explanation of this technique in the Appendix, a parameter k needs to be defined for this technique. Based on experience and several tests, a good value for k is the square root of the length of the input dataset. The following steps were taken to implement the k-Nearest Neighbor algorithm in R:

1. Load the data with the function "read.csv"
2. Normalize all predictors values with the function "scale"
3. Use the "class" library to create the k-NN model on the training dataset with the square root of the length of the dataset as the value of k
4. Build the Confusion Matrix for the prediction on validation se
5. Run the model on the testing dataset

6.2.1 Prediction of the issuance of volume-related reroute without specifying its type

Table 13 provides the confusion matrix for the first prediction on the testing set. The model had an accuracy of 0.5675, a Kappa statistic of 0.2523, an MCC of 0.3427, a sensitivity of 0.957 and a specificity of 0.3696.

Table 13: Confusion Matrix of k-NN predictions on the issuance of volume-related reroute advisories

	Actual No Reroute	**Actual Yes Reroute**
Predicted No Reroute	859	59
Predicted Yes Reroute	1414	1091

Table 13 shows that the model predicted correctly 859 No Reroute instances and 1091 Yes Reroute instances. However, it also predicted incorrectly 59 Yes Reroute and 1414 No Reroute. While the low specificity value shows that the model was limited in its ability to predict the No Reroute instances, the high sensitivity value shows that it was good at predicting Yes Reroute instances.

The overall performance of the k-NN technique for the first prediction is measured by both Kappa statistic (0.2523) and MCC (0.3427). Based on these values and on the interpretation of Landis and Koch [41], the k-NN model can be considered as a fair classifier for this prediction. Its low performance in predicting the No reroute instances penalizes it a lot since there are more instances of No Reroute.

6.2.2 Prediction of the issuance of volume-related reroute with specifying its type

Table 14 provides the confusion matrix for the prediction of the issuance of volume-related reroute advisories with specifying the reroute type. The model had an accuracy of 0.5998 and a Kappa statistic of 0.3062.

Table 14: Confusion Matrix of k-NN predictions on the issuance and the type of volume-related reroute advisories

	Actual No Reroute	**Actual RMD Reroute**	**Actual RQD Reroute**
Predicted No Reroute	1097	12	67
Predicted RMD Reroute	0	0	0
Predicted RQD Reroute	1146	129	932

Table 14 shows that the model correctly predicted No Reroute instances but incorrectly predicted 12 RMD Reroute advisories, and 67 RQD Reroute advisories. It did not predict any RMD Reroute class. Finally it predicted correctly 932 RQD Reroute advisories but incorrectly 129 RMD Reroute advisories, and 1146 No Reroute situations.

Table 15 summarizes the Sensitivity and Specificity metrics computed for all classes of the Decision Tree model. The moderate sensitivity of the No Reroute class shows that the model has not been able to predict the majority of No Reroute class. The moderate specificity (0.47) of the RQD Reroute class shows that less than half of the model's RQD Reroute prediction were correct. We can indeed see that 1146 of them were actually No Reroute situations.

Table 15: Sensitivity and Specificity metrics for each class of the prediction model with the k-NN technique

	Class: No Reroute	**Class: RMD Reroute**	**Class: RQD Reroute**
Sensitivity	0.4891	0	0.9329
Specificity	0.9307	1	0.4652

The overall performance of the Decision Tree algorithm for the second prediction is measured with the Kappa statistic of 0.3062. Based on this value and on the interpretation of Landis and Koch [41], the k-NN model can be considered as a fair classifier for the second prediction. The k-NN technique is known to performed poorly on unbalanced dataset because it is very sensitive to the scale of the data. The absence of prediction on the RMD Reroute class confirms this statement.

6.3 Naïve-Bayes

The following steps were taken to implement the Naïve-Bayes technique in R:

1. Load the data in R with the "read.csv" function

2. Create the model with the “naiveBayes” function from the R library “e1071” and train it the training dataset
3. Use the validation set and the “predict” function to test the model
4. Run the model on the testing dataset

6.3.1 *Prediction of the issuance of volume-related reroute without specifying its type*

Table 16 provides the confusion matrix for the first prediction on the testing set. The model had an accuracy of 0.8365, a Kappa statistic of 0.6369, a MCC of 0.637, a sensitivity of 0.7728 and a specificity of 0.8689.

Table 16: Confusion Matrix of Naive-Bayes predictions on the issuance of volume-related reroute advisories

	Actual No Reroute	**Actual Yes Reroute**
Predicted No Reroute	1949	259
Predicted Yes Reroute	294	881

Table 16 shows that the model predicted correctly 1949 No Reroute instances and 881 Yes Reroute instances. However, it also predicted incorrectly 259 Yes Reroute and 294 No Reroute. Since the specificity value is higher than the sensitivity one, this model was better at predicting the No Reroute situations than the Yes Reroute.

The overall performance of Naïve-Bayes technique for the first prediction is measured by both Kappa statistic (0.6369) and MCC (0.637). Based on these values and on the interpretation of Landis and Koch [41], the Naïve Bayes model can be considered as a substantial classifier for this prediction.

6.3.2 Prediction of the issuance of volume-related reroute with specifying its type

Table 17 provides the confusion matrix for the prediction of the issuance of volume-related reroute advisories with specifying the reroute type. The model had an accuracy of 0.8096 and a Kappa statistic of 0.5889.

Table 17: Confusion Matrix of Naïve-Bayes prediction on the issuance and the type of volume-related reroute advisories

	Actual No Reroute	**Actual RMD Reroute**	**Actual RQD Reroute**
Predicted No Reroute	1965	31	226
Predicted RMD Reroute	5	5	4
Predicted RQD Reroute	273	105	769

Table 17 shows that the model correctly predicted 1965 No Reroute instances but incorrectly predicted 31 RMD Reroute advisories, and 226 RQD Reroute advisories. It also predicted correctly 5 RMD Reroute advisories, and incorrectly 4 RQD Reroute advisories and 5 No Reroute instances. Finally, it predicted correctly 769 RQD Reroute advisories but incorrectly 105 RMD Reroute advisories, and 273 No Reroute instances.

Table 18 summarizes the Sensitivity and Specificity metrics computed for all classes of the Naïve-Bayes model. The very low value of the RMD Reroute class sensitivity shows that the model was not able to predict recommended reroute advisories. This pattern confirms the known weakness of Naïve Bayes technique: it relies on the assumption of equally important and independent features which is certainly not the case in this problem.

Table 18: Sensitivity and Specificity metrics for each class of the prediction model with Naïve-Bayes technique

	Class: No Reroute	Class: RMD Reroute	Class: RQD Reroute
Sensitivity	0.8761	0.035	0.7698
Specificity	0.7746	0.997	0.8414

The overall performance of the Naïve-Bayes algorithm for the second prediction is measured with the Kappa statistic of 0.5889. Based on this value and on the interpretation of Landis and Koch [41], the Naïve-Bayes model can be considered as a moderate classifier for the second prediction.

6.4 Support Vector Machines (SVM)

The following steps were taken to implement the SVM technique in R:

1. Use the "read.csv" function to import the dataset into R
2. Use the "ksvm" function to train the model on the training set
3. Similarly to Naïve-Bayes, use the validation dataset and the function "predict" to test the model
4. Create the confusion matrix for the prediction with the validation dataset
5. Run the model on the testing dataset and create the appropriate confusion matrix

6.4.1 Prediction of the issuance of volume-related reroute without specifying its type

Table 19 provides the confusion matrix for the first prediction on the testing set. The model had an accuracy of 0.8903, a Kappa statistic of 0.7561, an MCC of 0.7562, a sensitivity of 0.85 and a specificity of 0.9108.

Table 19: Confusion Matrix of SVM predictions on the issuance of volume-related reroute advisories

	Actual No Reroute	**Actual Yes Reroute**
Predicted No Reroute	2043	171
Predicted Yes Reroute	200	969

Table 19 shows that the model predicted correctly 2043 No Reroute instances and 969 Yes Reroute instances. However, it also predicted incorrectly 171 Yes Reroute and 200 No Reroute. Both sensitivity and specificity values are relatively high which indicate that the model generally performs well. The specificity value being a little higher than the sensitivity one, it can be concluded that this model was better at predicting the No Reroute instances than the Yes Reroute.

The overall performance of SVM technique for the first prediction is measured by both Kappa statistic (0.7561) and MCC (0.7562). Based on these values and on the interpretation of Landis and Koch [41], the SVM model can be considered as a substantial classifier for this prediction.

6.4.2 Prediction of the issuance of volume-related reroute with specifying its type

Table 20 provides the confusion matrix for the prediction of the issuance of volume-related reroute advisories with specifying the reroute type. The model had an accuracy of 0.8658 and a Kappa statistic of 0.7081.

Table 20: Confusion Matrix of SVM prediction on the issuance and the type of volume-related reroute advisories

	Actual No Reroute	Actual RMD Reroute	Actual RQD Reroute
Predicted No Reroute	2075	36	172
Predicted RMD Reroute	3	33	6
Predicted RQD Reroute	165	72	821

The Table 20 shows that the model correctly predicted 2075 No Reroute instances but incorrectly predicted 36 RMD Reroute advisories, and 172 RQD Reroute advisories. It also predicted correctly 33 RMD Reroute advisories, and incorrectly 6 RQD Reroute advisories and 3 No Reroute instances. Finally, it predicted correctly 821 RQD Reroute advisories but incorrectly 72 RMD Reroute advisories, and 165 No Reroute instances.

Table 21 summarizes the Sensitivity and Specificity metrics computed for all classes of the SVM model. All values are higher than those of the Naïve-Bayes model which confirms that the SVM model performs better. However, it is still remarkable that the sensitivity of the RMD Reroute class is much lower than that of the other classes. Thus, even if it performs well on predicting No Reroute and RQD reroute classes with few mistakes, the SVM model is still limited in its ability to correctly predict recommended reroute advisories.

Table 21: Sensitivity and Specificity metrics for each class of the prediction model with SVM technique

	Class: No Reroute	Class: RMD Reroute	Class: RQD Reroute
Sensitivity	0.9251	0.234	0.8218
Specificity	0.8175	0.9972	0.90

The overall performance of the SVM algorithm for the second prediction is measured with the Kappa statistic of 0.7081. Based on this value and on the interpretation of Landis and Koch [41], the SVM model can be considered as a substantial classifier for the second prediction.

6.5 Bagging Ensembles

The following steps were taken to implement the Bagging Ensembles technique in R:

1. Use the "read.csv" function to import the dataset into R
2. Use the "bagging" function to train the model with the training dataset
3. Use the validation dataset and the function "predict" to test the model
4. Create the confusion matrix for the prediction with the validation dataset
5. Run the model on the testing dataset and create the appropriate confusion matrix

6.5.1 Prediction of the issuance of volume-related reroute without specifying its type

Table 22 is the confusion matrix for the first prediction on the testing set. The model had an accuracy of 0.8563, a Kappa statistic of 0.6792, an MCC of 0.6792, a sensitivity of 0.7912 and a specificity of 0.8894.

Table 22: Confusion Matrix of Bagging Ensembles predictions on the issuance of volume-related reroute advisories

	Actual No Reroute	**Actual Yes Reroute**
Predicted No Reroute	1995	238
Predicted Yes Reroute	248	902

Table 22 shows that the model predicted correctly 1995 No Reroute instances and 902 Yes Reroute instances. However, it also predicted incorrectly 238 Yes Reroute and

248 No Reroute. Once again, the specificity value is a higher than the sensitivity one. Thus, this model was better at predicting the No Reroute instances than it was at predicting the Yes Reroute ones.

The overall performance of Bagging Ensembles technique for the first prediction is measured by both Kappa statistic (0.6792) and MCC (0.6792). Based on these values and on the interpretation of Landis and Koch [41], the Bagging Ensembles model can be considered as a substantial classifier for this prediction.

6.5.2 Prediction of the issuance of volume-related reroute with specifying its type

Table 23 provides the confusion matrix for the prediction of the issuance of volume-related reroute advisories with specifying the reroute type. The model had an accuracy of 0.9199 and a Kappa statistic of 0.8295.

Table 23: Confusion Matrix of Bagging Ensembles prediction on the issuance and the type of volume-related reroute advisories

	Actual No Reroute	**Actual RMD Reroute**	**Actual RQD Reroute**
Predicted No Reroute	2127	15	98
Predicted RMD Reroute	8	96	12
Predicted RQD Reroute	108	30	889

Table 23 shows that the model correctly predicted 2127 No Reroute instances but incorrectly predicted 15 RMD Reroute advisories, and 98 RQD Reroute advisories. It also predicted correctly 96 RMD Reroute advisories, and incorrectly 12 RQD Reroute advisories and 8 No Reroute instances. Finally, it predicted correctly 889 RQD Reroute advisories but incorrectly 30 RMD Reroute advisories, and 108 No Reroute instances.

Table 24 summarizes the Sensitivity and Specificity metrics computed for all classes of the Bagging Ensembles model. While the sensitivity value of the RMD Reroute class was low for the last three techniques, the value obtained for Bagging Ensembles algorithm is in pair with that of the Decision Tree algorithm. It means that this model performs much better on this class (RMD Reroute) while also performing well on the other two classes.

Table 24: Sensitivity and Specificity metrics for each class of the prediction model with Bagging Ensembles technique

	Class: No Reroute	Class: RMD Reroute	Class: RQD Reroute
Sensitivity	0.9483	0.6809	0.8899
Specificity	0.9009	0.9938	0.9421

The overall performance of the Bagging Ensembles algorithm for the second prediction is measured with the Kappa statistic of 0.8295. Based on this value and on the interpretation of Landis and Koch [41], the Bagging Ensembles model can be considered as an almost perfect classifier for the second prediction. This model's performance metrics and those of SVM model have common trends (high specificities and unbalanced sensitivities). However, its higher sensitivities values for RMD Reroute and RQD Reroute classes makes this model more accurate on the less frequent classes. Its Kappa statistic is thus higher and also superior to the "almost perfect classifier" threshold value 0.81.

6.6 Boosting Ensembles

The following steps were taken to implement the Boosting Ensembles technique in R:

1. Use the "read.csv" function to import the dataset into R
2. Use the "boosting" function to train the model with the training dataset
3. Use the validation dataset and the function "predict" to test the model

4. Create the confusion matrix for the prediction with the validation dataset
5. Run the model on the testing dataset and create the appropriate confusion matrix

6.6.1 Prediction of the issuance of volume-related reroute without specifying its type

Table 25 provides the confusion matrix for the first prediction on the testing set. The model had an accuracy of 0.9172, a Kappa statistic of 0.816, an MCC of 0.8161, a sensitivity of 0.8904 and a specificity of 0.9309.

Table 25: Confusion Matrix of Boosting Ensembles predictions on the issuance of volume-related reroute advisories

	Actual No Reroute	**Actual Yes Reroute**
Predicted No Reroute	2088	125
Predicted Yes Reroute	155	1015

Table 25 shows that the model predicted correctly 2088 No Reroute instances and 1015 Yes Reroute instances. However, it also predicted incorrectly 125 Yes Reroute and 155 No Reroute. Both sensitivity and specificity are high for a classification problem but the fact that the sensitivity is higher than the specificity indicates that the model is better at predicting the No Reroute class than the Yes Reroute one. Those performance metrics share similar patterns with the Decision Tree ones.

The overall performance of the Boosting Ensembles technique for the first prediction is measured by both Kappa statistic (0.816) and MCC (0.8162). Based on these values and on the interpretation of Landis and Koch [41], the Boosting Ensembles model can be considered as an almost perfect classifier for this prediction.

6.6.2 *Prediction of the issuance of volume-related reroute with specifying its type*

Table 26 provides the confusion matrix for the prediction of the issuance of volume-related reroute advisories with specifying the reroute type. The model had an accuracy of 0.8685 and a Kappa statistic of 0.7154.

Table 26: Confusion Matrix of Boosting Ensembles prediction on the issuance and the type of volume-related reroute advisories

	Actual No Reroute	**Actual RMD Reroute**	**Actual RQD Reroute**
Predicted No Reroute	2072	31	157
Predicted RMD Reroute	7	29	5
Predicted RQD Reroute	164	81	837

Table 26 shows that the model correctly predicted 2072 No Reroute instances but incorrectly predicted 31 RMD Reroute advisories, and 157 RQD Reroute advisories. It also predicted correctly 29 RMD Reroute advisories, and incorrectly 5 RQD Reroute advisories and 7 No Reroute instances. Finally, it predicted correctly 837 RQD Reroute advisories but incorrectly 81 RMD Reroute advisories, and 164 No Reroute instances.

Table 27 summarizes the Sensitivity and Specificity metrics computed for all classes of the Boosting Ensembles model. The sensitivity and specificity values share common patterns with the SVM model with a very low sensitivity on the RMD Reroute class, but also worse prediction capabilities on the No Reroute and RQD Reroute classes than the Bagging Ensembles model. The model is therefore decent to predict No Reroute and RQD Reroute classes but is limited in its ability to predict RMD Reroute class.

Table 27: Sensitivity and Specificity metrics for each class of the prediction model with Boosting Ensembles technique

	Class: No Reroute	Class: RMD Reroute	Class: RQD Reroute
Sensitivity	0.9238	0.2057	0.8378
Specificity	0.8351	0.9963	0.8972

The overall performance of the Boosting Ensembles algorithm for the second prediction is measured with the Kappa statistic of 0.7154. Based on this value and on the interpretation of Landis and Koch [41], the Boosting Ensembles model can be considered as a substantial classifier for the second prediction. It is once again penalized by its performance on the RMD Reroute class which shows that this algorithm is limited when predicting the classes with low number of instances.

6.7 Random Forest

The following steps were taken to implement the Random Forest technique in R:

1. Use the “read.csv” function to import the dataset into R
2. Use the “randomForest” function to train the model with the training dataset
3. Use the validation dataset and the function “predict” to test the model
4. Create the confusion matrix for the prediction with the validation dataset
5. Run the model on the testing dataset and create the appropriate confusion matrix

6.7.1 Prediction of the issuance of volume-related reroute without specifying its type

Table 28 provides the confusion matrix for the first prediction on the testing set. The model had an accuracy of 0.9279, a Kappa statistic of 0.84, an MCC of 0.84, a sensitivity of 0.9096 and a specificity of 0.9371.

Table 28: Confusion Matrix of Random Forest predictions on the issuance of volume-related reroute advisories

	Actual No Reroute	Actual Yes Reroute
Predicted No Reroute	2102	103
Predicted Yes Reroute	141	1037

The Table 28 shows that the model predicted correctly 2102 No Reroute instances and 1037 Yes Reroute instances. However, it also predicted incorrectly 103 Yes Reroute and 141 No Reroute. The sensitivity and specificity values are the highest obtained among all techniques considered. The high sensitivity in particular shows that this technique performed very well on the Yes Reroute instances which was in general not the case with other techniques.

The overall performance of the Random Forest technique for the first prediction is measured by both Kappa statistic (0.84) and MCC (0.84). Based on these values and on the interpretation of Landis and Koch [41], the Random Forest model can be considered as an almost perfect classifier for this prediction.

6.7.2 Prediction of the issuance of volume-related reroute with specifying its type

Table 29 provides the confusion matrix for the prediction of the issuance of volume-related reroute advisories with specifying the reroute type. The model had an accuracy of 0.906 and a Kappa statistic of 0.7984.

Table 29: Confusion Matrix of Random Forest prediction on the issuance and the type of volume-related reroute advisories

	Actual No Reroute	Actual RMD Reroute	Actual RQD Reroute
Predicted No Reroute	2122	22	113
Predicted RMD Reroute	6	72	15
Predicted RQD Reroute	115	47	871

Table 29 shows that the model correctly predicted 2122 No Reroute instances but incorrectly predicted 22 RMD Reroute advisories, and 113 RQD Reroute advisories. It also predicted correctly 72 RMD Reroute advisories, and incorrectly 15 RQD Reroute advisories and 6 No Reroute instances. Finally, it predicted correctly 871 RQD Reroute advisories but incorrectly 47 RMD Reroute advisories, and 115 No Reroute instances.

Table 30 summarizes the Sensitivity and Specificity metrics computed for all classes of the Random Forest model. Both sensitivity and specificity values are similar to those of the Bagging Ensembles model. Thus, the model predicts decently the RMD Reroute class and well the No Reroute and RQD Reroute classes.

Table 30: Sensitivity and Specificity metrics for each class of the prediction model with Random Forest technique

	Class: No Reroute	Class: RMD Reroute	Class: RQD Reroute
Sensitivity	0.9461	0.5106	0.8719
Specificity	0.8816	0.9935	0.932

The overall performance of the Random Forest algorithm for the second prediction is measured with the Kappa statistic of 0.7984. Based on this value and on the interpretation of Landis and Koch [41], the Random Forest model can be considered as a substantial classifier for the second prediction but close from the almost perfect classifier. It appears to be performing better than the Boosting Ensembles model on low instances classes but cannot reach the performance of the Decision Tree and Bagging Ensembles models.

6.8 Comparison of techniques

6.8.1 Prediction of the issuance of volume-related reroute advisories

For this problem, the prediction model was predicting two classes:

- No Reroute: no volume-related reroute advisory is being issued
- Yes Reroute: at least one volume-related reroute advisory is being issued

Since only two prediction classes are possible for this problem, the performance of the prediction model can be measured with Matthew's Coefficient, also called MCC. Figure 25 displays the MCC value for both validation and testing datasets for each technique.

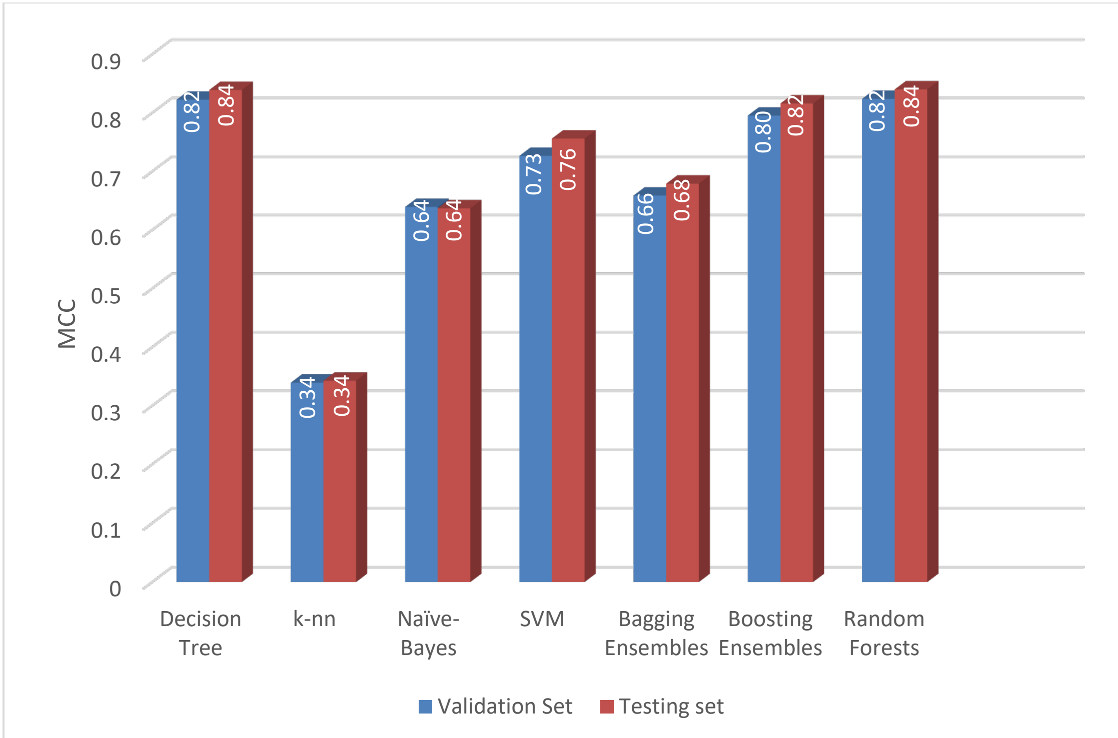

Figure 25: Comparison of various ML techniques for predicting the issuance of volume-related reroute advisory using Matthew's coefficient (MCC) for validation and testing datasets

According to Figure 25, the **Random Forest** technique appears to be the best performing one based on MCC value, closely followed by Decision Tree. Thus, within all seven Machine Learning that have been tested, **Random Forest** algorithm was the best suited for predicting the issuance of volume-related algorithms. It is noticeable that Decision Tree technique obtained very close results and could also be chosen.

6.8.2 Prediction of the issuance and the type of volume-related reroute advisories

To predict the issuance and the type of the reroute advisory, three prediction classes were considered:

– No Reroute: no volume-related reroute advisory is being issued

– RQD Reroute: required reroute issued by Air Traffic Controllers

– RMD Reroute: recommended reroute issued by Air Traffic Controllers

Because this model can predict three classes, it is not possible to use Matthew's Coefficient anymore to measure the performances of each technique. The Kappa statistic appears to be the most reliable performance measure because it is very appropriate for evaluating unbalanced datasets. Figure 26 provides the Kappa statistic for both validation and testing datasets for each technique.

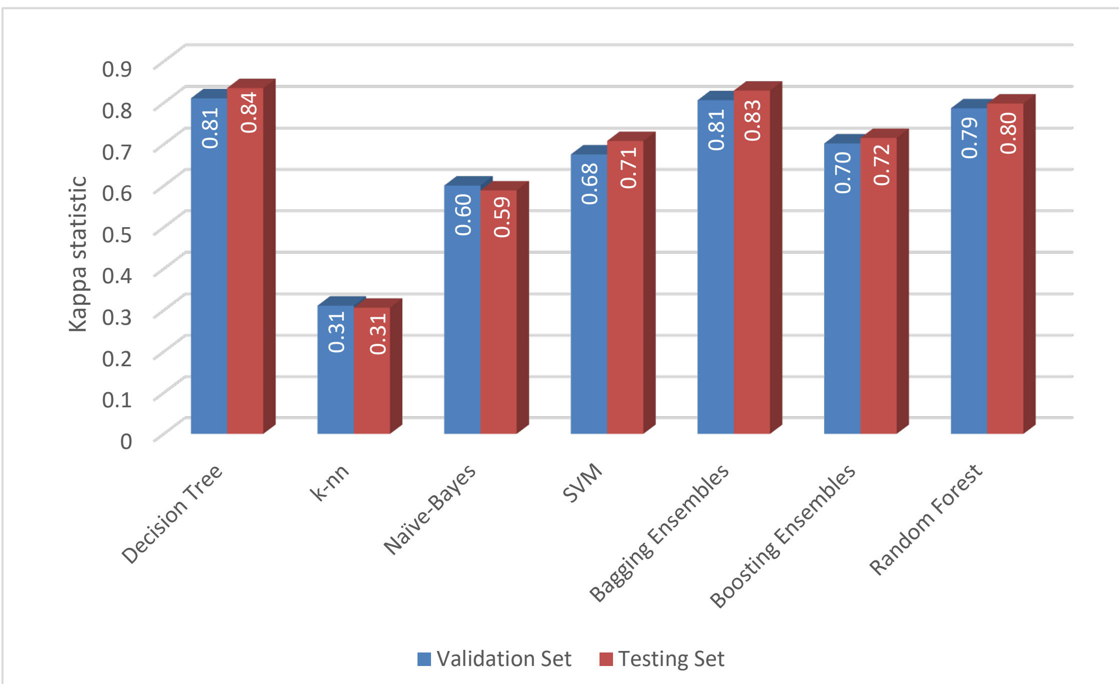

Figure 26: Comparison of various ML techniques for predicting the issuance and the type of volume-related reroute advisory using Kappa statistic

From Figure 26, it appears that the **Decision Tree** technique was the best performing one among all seven techniques that have been tested. The Bagging Ensembles technique reached similar level of performance but was penalized by its lower performance on the RMD Reroute class.

6.9 Chapter Conclusion

Seven Machine Learning techniques known for their classification abilities (Decision Tree, Nearest Neighbor, Naïve-Bayes, Support Vector Machines, Bagging Ensembles, Boosting Ensembles, and Random Forest) were benchmarked in order to identify the best performing model to predict the issuance and the type of volume-related reroute advisories. The dataset considered consists of a large database that includes all reroutes issued for four months (January - April 2017) fused with the traffic count per facility of the national airspace. This traffic count was self-generated by the author because the processing time to extract the actual would have taken longer than expected. Thus, the work conducted in this chapter should be considered as a proof of concept since the dataset has been generated "manually" by the author for demonstration purposes. The Machine Learning techniques were them tested on this dataset to provide two predictions:

- Prediction of the issuance of volume-related reroute advisories without specifying the reroute type
- Prediction of the issuance of volume-related advisories with specifying the reroute type (Required, Recommended)

For the first prediction, the model was evaluated using Matthew's Coefficient, while for the second one, the Kappa statistic was used. For the first prediction, the Random Forest technique appeared to be the best performing one with a MCC of 0.84, followed very closely by the Decision Tree technique with a MCC of 0.839. The Decision Tree technique also appeared to be the best suited for the second prediction with a Kappa statistic of 0.8356.

The work presented in this chapter is therefore a proof of concept that a prediction model can be developed with a classification technique such as Decision Tree to predict the issuance and the type of volume-related reroute advisories. If the hourly traffic count per facility can be extracted by the FAA, it will then be possible to run the seven Machine Learning techniques in order to have a realistic prediction model. This work thus provides an answer to **Research Question #2** and validates **Hypothesis #2.**

CHAPTER 7. CONCLUSION AND FUTURE WORK

7.1 Conclusion

Reroute advisories are Traffic Management Initiatives issued when an Air Route Traffic Control Center (ARTCC) identifies constraint(s) and assigns new routes to affected flights. They are defined by their level of urgency that can be either Required, Recommended or For Your Information (FYI). While most of the past research conducted on reroute advisories focused on the definition and the optimization of reroutes, none of the past research analyzed the relevancy of reroute advisories and considered their urgency levels (Required, Recommended, FYI). Furthermore, a research gap has been identified on the collection of relevant datasets to build prediction models in order to predict the issuance of reroute advisories. This gap is partially fulfilled in this work in which a model has been created to predict the issuance and the type of volume-related reroute advisories.

7.1.1 Assessment of the relevancy of recommended reroutes

While pilots almost always comply with required reroutes, their decisions regarding recommended reroutes vary. There is thus a need for the Federal Aviation Administration (FAA) to assess the efficiency and relevance of Recommended Reroutes to identify optimal routes for future events and to have a better understanding of pilots' decisions. The first **Research Question** is directly related to that need and the hypothesis formulated to answer the research question suggests a solution that have been implemented by the author.

- **Research Question 1:** How can the relevancy of recommended reroutes be best captured?

- **Hypothesis 1:** If Data Fusion is used to analyze recommended reroute messages and flight data, then trends and patterns may be identified and used in assessing the relevancy of recommended reroutes.

A four-steps methodology has been developed and implemented to answer the Research Question 1. It has been applied on four months of data provided to the author by the FAA. First, System Wide Information Management (SWIM) Flight Data Publication Service (SFDPS) and Traffic Flow Management System (TFMS) datasets have been identified to provide respectively, flight information (tracking, flight plans, etc.), and traffic data issued by Air Traffic Controllers such as Reroute advisories. Data has then been processed to extract only relevant information for the scope of this Research Question and find correlations between datasets such as airports affected by recommended reroutes and origin/destination airports. Those correlations have been used to fuse the SFDPS and TFMS datasets in order to extract all flights affected by recommended reroutes during the period considered (January - April 2017). The author finally developed an algorithm and metrics to compute flights compliance to reroutes and assess the relevancy of recommended reroutes during the aforementioned period.

According to the compliance metric selected, it appeared that only 2.1% of all flights affected by reroutes chose their recommended reroute. This work also highlighted that recommended reroutes affects mostly the East Coast airspace and seem to be designed for mid-distance flights (1550 – 2500 km). These results and those presented in Section 5.5 provide Air traffic Controllers with information as to the relevancy of reroutes options and therefore may help them to issue more relevant reroutes in the future.

Finally, the FAA may benefit from the Python scripts developed within the scope of this research. As mentioned, Python scripts have been developed to:

- Assess the compliance of pilots to reroute advisories over a certain period of time
- Assess the closeness of two routes based on geographical positions

These two scripts may be re-used by the FAA for purposes close to this research. It would for example be interesting to change the analyzed time window and assess the impact of the time window on reroutes advisories. This work could indeed be scaled to a larger analysis (e.g. one full year) by just integrating more data and allocating more computational time.

The SFDPS and TFMS parsers developed may also prove useful to the FAA because these datasets are involved in many different research fields.

7.1.2 Prediction of volume-related reroute advisories

Reroute advisories can be issued for three main reasons: weather conditions, volume constrains or equipment issues. Because the description of bad weather conditions is not detailed in reroute advisories, it is therefore very complex to build a database matching all weather-related reroute advisories with the corresponding weather situation. Equipment issues are hard to forecast and lead to the issuance of reroute advisories much less frequently than weather and volume constraints. According to these assumptions, the author focused on volume-related advisories and the following Research Question and Hypothesis were formulated:

- **Research Question 2:** How can the issuance of a volume-related reroute advisory be accurately predicted?
- **Hypothesis 2:** If traffic data is fused with reroute data (TFMS), and supervised machine learning algorithms are used to develop prediction models, then it will be possible to find the algorithm that best predicts the issuance of volume-related reroute advisories.

Seven Machine Learning (Decision Tree, Nearest Neighbor, Naïve-Bayes, Support Vector Machines, Bagging Ensembles, Boosting Ensembles, and Random Forest) techniques, known for their classification ability, were benchmarked in order to identify the model that best predicts the issuance and the type of volume-related reroute advisories. The dataset consists of a large database that includes all reroutes issued for four months (January - April 2017) fused with the traffic count per facility of the national airspace. This traffic count was manually generated by the author because the processing time to extract this data would have been taken longer than expected. Thus, the work conducted in this chapter should be considered as a proof of concept since the dataset was generated for demonstration purposes only. The Machine Learning techniques were tested on the dataset to provide two predictions:

- Prediction of the issuance of volume-related reroute advisories without specifying the reroute type
- Prediction of the issuance of volume-related advisories with specifying the reroute type (Required, Recommended)

For the first prediction, the models were evaluated using Matthew's Coefficient, while for the second one, the Kappa statistic was used. For the first prediction, the Random Forest technique appeared to be the best performing one with a MCC of 0.84, followed very closely by the Decision Tree technique with a MCC of 0.839. The Decision Tree technique also appeared to be the best suited for the second prediction with a Kappa statistic of 0.8356.

7.2 Future work

7.2.1 Assessment of the relevancy of recommended reroutes

This work detailed a first approach to assess the relevancy of recommended flight reroutes. Within the context of this research, four metrics were developed to assess the compliance of flights to the recommended reroute they were affected by. Eventually one was identified as the best performing one. Future work would consider defining an additional metric based on the lateral deviation of the flight from its reroute. This metric would provide another point of view because it would not rely on any of the two approaches (polygons, circles) presented in this work. Furthermore, if different accurate metrics are defined, it may be relevant to mathematically combine them so as to develop a general compliance metric aggregating different approaches together.

Moreover, this work focused on a 2-D (x-y) analysis of recommended reroutes. However, reroutes advisories most of the time provide altitude boundaries for flights such that reroutes are defined in three dimensions. It could be possible to incorporate the altitude of the plane into this work because this information is accessible in the SFDPS dataset. The main problem is that the altitude requirement is usually hand-written by Air Traffic

Controllers in the remark section and therefore a more complex algorithm would be needed to always understand such requirement.

Finally, this work focused mainly on traffic data provided by the Federal Aviation Administration but did not take into account other parameters that interfere with pilots' decisions and that are related to the plane itself. For example, pilots' choices are certainly influenced by economic considerations such as fuel efficiency/consumption, time considerations such as the potential delays, etc. Consequently, it would be very interesting to acquire information about those parameters and incorporate them into this study. Doing so would help better informs the reasons for pilot's decisions to not follow a recommended reroute.

7.2.2 Prediction of volume-related reroute advisories

The work presented in this research to predict the issuance of volume-related reroute advisories is a proof of concept since the traffic count data per facility has been generated by the author. Future work would focus on collecting the actual hourly traffic count data for a long period of time and re-train the Machine Learning models. This would help indicate whether the performance of each algorithm would hold and whether the Decision Tree technique would still be the best technique to predict the issuance of volume-related reroute advisories.

Finally, because weather-related reroute advisories are the most frequent one, future work could also focus on predicting their issuance. As explained earlier, the main challenge that would need to be addressed is that reroute advisories collected in TFMS do not provide enough detail as the weather to the conditions causing their issuance.

However, this gap may be filled by working closely with the FAA on this subject to identify relevant datasets that would contain the needed information.

APPENDIX A: MACHINE LEARNING TECHNIQUES

A.1 Decision Trees

Decision Trees is a classifying Machine Learning technique that relies on tree-structure in order to sort instances based on their feature values. Nodes of the tree represent the features while branches represent the instance's value for the corresponding feature. The root node of the tree corresponds to the feature that divides the most comprehensively the training dataset. Each subset is then always divided according to the feature that best divides it. Repeating this process and thus creating sub-trees leads to breaking down the training dataset into same classes subsets. [43] Once the tree is complete, it is then possible for the classifier to predict new instances.

A.2 Support Vector Machines

The Support Vector Machine (SVM) technique was initially developed for statistical learning theory and later adopted in Machine Learning theory and signal processing [44]. SVM algorithms are very appropriate for solving classification problems and pattern recognition problems such as object recognition or speaker identification. The SVM technique is based on the notion of "margin", the two sides of a hyperplane that splits two data classes. "Maximizing the margin and thereby generating largest possible distance between the separating hyperplane and the data on either side of it has been proven to reduce the expected generalisation error" [45]. In other words, the algorithm should be provided with a labelled training data in order to output an optimal hyperplane which classifies new examples.

A.3 k-Nearest Neighbour (kNN)

The kNN technique is part of the instance-based learning algorithms. It relies on the observation that instances in a dataset will most of the time have strong similarities with other instances of the same dataset. Thus the label of an unclassified instance will be determined by looking at the class of its nearest neighbours and identifying the most frequent class label within the neighbours. The parameter k, that needs to be set by the user, actually corresponds to the number of neighbours the algorithm is observing for each instance. It is commonly known that the square root of the dataset's size is a good assumption for k. [43]

A.4 Naïve Bayes

The Naïve Bayes classification technique is one of the statistical learning algorithms. It relies on the calculation of probabilities of each label based on the observation of the classified instances of the training dataset. Unclassified instances are then predicted using the most likely class for all features. Naïve Bayes classifier is known for the short computational time its training phase requires.

A.5 Bagging and Boosting Ensembles

These two techniques are members of the "meta-algorithms" family of Machine learning techniques. These techniques combine different machine learning algorithms in order to decrease precise features (e.g. variance, bias).

Bagging Ensembles stands for Bootstrap Aggregating Ensembles and its main objective is to decrease the variance of the prediction. This objective is fulfilled by increasing the size of the training dataset using combinations and repetitions.

Boosting Ensembles first uses the original training dataset to produce models that perform averagely. These models are then boosted by creating new subsets containing the instances that were most of the time incorrectly classified with previous models.

A.6 Random Forest

Random Forest is an ensemble learning algorithm. It can be thought of as a collection of many decision trees and uses bootstrapping, like Bagging Ensembles method, to resample data from the training set. Random Forest algorithms are very efficient on large datasets and are known to avoid overfitting by themselves.

APPENDIX B: CODE

The code developed by the author for this Master Thesis and used to obtain the results presented in this document is available on GitHub platform and can be found at the following address: https://github.com/temanava/Ms-Thesis-Reroute-Advisories.git

This code is not exhaustive but gathers the key steps taken by the author during his work. It relies on the acquisition of SFDPS and TFMS datasets used in this work. These datasets are not public and have been provided by the FAA to the author. Thus, the code provided in this GitHub repository is made public by the author for indicative purposes but cannot be run without the data.

REFERENCES

[1] "National Airspace System Overview," [Online]. Available: https://www.faa.gov/air_traffic/nas/nynjphl_redesign/documentation/feis/media/Appendix_A-National_Airspace_System_Overview.pdf.

[2] "Air Traffic Control System Command Center (ATCSCC)," 25 February 2014. [Online]. Available: https://www.faa.gov/about/office_org/headquarters_offices/ato/service_units/systemops/nas_ops/atcscc/.

[3] M. Vanhoenacker, "In Flight, Image from Navtech," 14 May 2015. [Online]. Available: https://www.nytimes.com/interactive/2015/05/14/opinion/14-in-flight-mark-vanhoenacker.html, image from Navtech.

[4] FAA Seal, "Air Traffic By The Numbers," 14 November 2017. [Online]. Available: www.faa.gov/air_traffic/by_the_numbers.

[5] NBAA Air Traffic Services (ATS), "Traffic Management Initiatives," [Online]. Available: https://filesmart.org/resources/traffic-management-initiatives/.

[6] National Business Aviation Association (NBAA), "Tools used for Traffic Flow Management - Reroutes," [Online]. Available: https://nbaa.org/aircraft-operations/airspace/tfm/tools-used-for-traffic-flow-management/reroutes/.

[7] National Business Aviation Association, "Tools used for Traffic Flow Management," [Online]. Available: https://www.nbaa.org/ops/airspace/tfm/tools/#terminal-tmi.

[8] C. Chuang, D. J. Lovell, M. O. Ball and G. Zhang, "An analysis of the relationship between Traffic Management Initiatives," ProQuest Dissertations and Theses, 2017.

[9] SkyVector, "SkyVector: Flight Planning / Aeronautical Charts," [Online]. Available: https://skyvector.com.

[10] Federal Aviation Administration, "Electronic Code of Federal Regulations," 19 October 2018. [Online]. Available: https://www.ecfr.gov/cgi-bin/text-idx?c=ecfr&sid=3efaad1b0a259d4e48f1150a34d1aa77&rgn=div5&view=text&node=14:2.0.1.3.10&idno=14#se14.2.91_1123.

[11] D. Bradbury, "How Big Data in Aviation Is Transforming the Industry," 4 March 2018. [Online]. Available: https://fr.hortonworks.com/article/how-big-data-in-aviation-is-transforming-the-industry/.

[12] A. Hassanian, A. Azar, A. Snášel, V. Kacprzyk, J. Abawajy, V. Snášel, Václav, Abawajy and H. Jemal , "Studies in Big Data, Volume 9," in *Big data in complex systems challenges and opportunities*, 2015.

[13] A. Samuel, "Some Studies in Machine Learning Using the Game of Checkers," *IBM Journal of Research and Development,* vol. 3, pp. 210-229, 1959.

[14] N. M. Wernick, Y. Yang, J. G. Brankov, G. Yourganov and S. C. Strother, "Machine Learning in Medical Imaging," *IEEE Signal Processing Magazine,* vol. 27, pp. 25-38, 2010.

[15] J. Rzeszotko and S. H. Nguyen, "Machine Learning for Traffic Prediction," *Fundamenta Informaticae,* vol. 119, no. 3-4, pp. 407-420, 2012.

[16] C. Turakhia, "Engineering More Reliable Transportationwith Machine Learning and AI at Uber," 10 November 2017. [Online]. Available: https://eng.uber.com/machine-learning/.

[17] J. O. Awoyemi, A. O. Adetunmbi and S. A. Oluwadare, "Credit card fraud detection using machine learning techniques: a comparative analysis," in *International Conference on Computing Networking and Informatics (ICCNI)*, Lagos, 2017.

[18] S. Suresh, N. Sundararajan and R. Savitha, Supervised learning with complex-valued neural networks (Studies in Computational Intelligence; 421), Berlin: Springer, Berlin, Heidelberg, 2013.

[19] K. Wakefield, "A guide to machine learning algorithms and their applications," [Online]. Available: https://www.sas.com/en_gb/insights/articles/analytics/machine-learning-algorithms.html.

[20] R. Sutton and A. Barto, Reinforcement Learning: An Introduction, Boston: MIT Press, 1998.

[21] L. a. T. Fellman, "ARTCC Initiated Rerouting," in *AIAA Aviation Technology*, 2006.

[22] P. Dutta, R. Bhattacharya and S. Chakravorty, "Robust Air traffic Control Using Ground Delays and Rerouting of Flights," *Aerospace Computing , Information and Communication,* vol. 6, no. 12, pp. 649-674, 2009.

[23] C. Gooding, A. Shelley, S. Wu, W. Johnson and C. Duong, "Pilot Convective Weather Decision Making in En Route Airspace," in *ICAS 2012 - 28th Congress of the International Council of the aeronautical Sciences*, Brisbane, Australia, 2012.

[24] J. Rios, "Aggregate Statistics of National Traffic Management Initiatives," in *10th AIAA Aviation Technology , Integration and Operations (ATIO) Conference*, Forth Worth, Texas, 2010.

[25] N. Nelson and M. Paglione, "Determination of Lateral and Vertical Adherence to Route," in *AIAA Guidance, Navigation, and Control (GNC) Conference*, Boston, 2013.

[26] H. S. Piquet, "Evaluation of new enroute performance measures for air navigation service providers," Atlanta, 2017.

[27] National Transportation Systems Center, "Swim Flight Data Publication Service (SFDPS) reference Manual," 2015.

[28] Federal Aviation Administration, "Traffic Flow Management Operational Support and Maintenance Statement of Work (SOW)," March 25, 2010.

[29] Federal Aviation Administration, "Java Messaging Service Description Document Traffic Flow Management Data Service (TFMData)," Federal avaiation Administration, 2016.

[30] L. Sorber, M. Van Barel and L. De Lathauwer, "Structured Data Fusion," *IEEE Journal of Selected Topics in Signal processing,* vol. 9, no. 4, pp. 586-600, June 2015.

[31] U.S. Department of Transportation, "28 Day NASR Subscription," Federal Aviation Administration, [Online]. Available: https://www.faa.gov/air_traffic/flight_info/aeronav/Aero_Data/NASR_Subscription/. [Accessed 2 February 2019].

[32] H. Lepori, "Introduction to FIXM," 2017. [Online]. Available: https://www.icao.int/MID/Documents/2017/SWIM.

[33] E. Mangortey, J. Gilleron, G. Dard, O. Pinon and D. N. Mavris, "Development of a Data Fusion Framework to support the Analysis of Aviation Big Data," in *AIAA Science and Technology Forum (AIAA Scitech)*, San Diego, 2019.

[34] P. Documentation, "19.7. xml.etree.elementtree - the elementtree xml api," [Online]. Available: docs.python.org/2/library/xml.etree.elementtree.html..

[35] Bokeh Development Team, "Bokeh: Python library for interactive visualization," [Online]. Available: https://bokeh.pydata.org/en/latest/.

[36] Federal Aviation Administration, "Air Traffic Activity System (ATADS)," [Online]. Available: https://aspm.faa.gov/opsnet/sys/main.asp. [Accessed 30 Marcg 2019].

[37] Office of Performance Analysis (AJR-G), Air Traffic Organization, Federal Aviation Administration, "Traffic Flow Management System," 2018.

[38] Sokolova and Lapalme, "A systematic analysis of performance measures for classification tasks," *Information Processing and Management,* vol. 45, no. 4, pp. 427-437, 2009.

[39] M. Sunasra, "Performance Metrics for Classification problems in Machine Learning," 11 November 2017. [Online]. Available: https://medium.com/greyatom/performance-metrics-for-classification-problems-in-machine-learning-part-i-b085d432082b.

[40] B. Lantz, Machine Learning with R: Discover How to Build Machine Learning Algorithms, Prepare Data, and Dig Deep into Data Prediction Techniques with R. Packt Publishing, Packt, 2015.

[41] J. R. Landis and G. G. Koch, "The Measurement of Observer Agreement for Categorical Data," *Biometrics,* vol. 33, pp. 159 - 174, 1977.

[42] Python Software Foundation, "13.1. csv - CSV File Reading and Writing," Python Software Foundation, [Online]. Available: https://docs.python.org/2/library/csv.html. [Accessed 25 February 2019].

[43] S. B. Kotsiantis, "Supervised Machine Learning: A review of Classification Techniques," Univerity of Peloponnese, Tripolis, 2007.

[44] Y. Ma and G. Guo, Support vector machines applications, 6 ed., 2014.

[45] F. Y. Osisanwo, J. E. Akinsola, O. Awodele, J. O. Hinmikaiye, O. Olakanmi and J. Akinjobi, "Supervised Machine Learning Algorithms: Classification and Comparison," *International Journal of Computer Trends and Technology (IJCTT),* vol. 48, no. 3, pp. 128 - 138, June 2017.

[46] "Image from National Air Traffic Controllers Association website," [Online]. Available: https://www.natca.org/.

[47] E. Minka, "Mean Absolute Error - MAE [Machine Learning (ML)]," 02 February 2018. [Online]. Available: https://medium.com/@ewuramaminka/mean-absolute-error-mae-machine-learning-ml-b9b4afc63077.

[48] S. Patel, "Chapter 2: SVM (support Vector Machine) - Theory - Machine Learning 101- Medium," 3 May 2017. [Online]. Available: https://medium.com/machine-learning-101/chapter-2-svm-support-vector-machine-theory-f0812effc72.

[49] R. Rew, G. Davis, S. Emmerson, H. Davies, E. Hartnett, D. Heimbigner and W. Fisher, "What is NetCDF," Unidata, [Online]. Available: https://www.unidata.ucar.edu/software/netcdf/docs/index.html.