A unified, clustering-based framework for detection of spatial and energy anomalies in trajectories utilizing ADS-B data

Samantha J. Corrado[1], Tejas G. Puranik[2], Olivia Pinon Fischer[3], and Dimitri N. Mavris[4]

Georgia Institute of Technology, Atlanta, GA 30332, United States

Abstract

As air traffic demand grows, robust, data-driven anomaly detection methods are required to ensure that aviation systems become safer and more efficient. The terminal airspace is identified as the most critical airspace for both individual flight-level and system-level safety and efficiency. As such, developing data-driven anomaly detection methods to analyze terminal airspace operations is paramount. With the expansion of ADS-B technology, open-source flight tracking data has become more readily available to enable larger-scale analyses of aircraft operations. This paper makes a distinction between spatial metrics in ADS-B trajectory data and energy metrics derived from ADS-B trajectory data. Motivated by the limited number of approaches that simultaneously consider both spatial and energy metrics, this paper introduces the concepts of spatial anomalies and energy anomalies. In particular, it proposes a novel, unified framework for detection of spatial and energy anomalies in ADS-B trajectory data (and associated derived metrics). The framework consists of three main parts - a data processing procedure, a spatial anomaly detection method, and an energy anomaly detection method. The framework is demonstrated utilizing four months of ADS-B trajectory data associated with arrivals at San Francisco International Airport, and the relationship between the spatial and energy anomalies in this terminal airspace is explored. The results that stem from the implementation of this framework indicate that if an aircraft is spatially not conforming to an identified set of air traffic flows representing standard spatial operations, then this aircraft is more likely to experience non-conformance to standard operations in its energy metrics. Aviation operators, such as air traffic controllers, may benefit from this observation, as it may factor into decision-making in instances where there is the potential to instruct an aircraft to spatially deviate from standard operations. Additionally, this research revealed underlying differences between trajectories that are spatially nominal yet energy-anomalous and those trajectories that are spatially anomalous and energy-anomalous. Focusing solely on energy anomaly detection does not provide insight into potential spatial-related decisions that may have been made to result in off-nominal energy behavior.

Keywords: air transportation, machine learning, anomaly detection, ADS-B, clustering

1. Introduction

In recent years, the aviation industry has seen a large increase in the volume of operations. According to the Federal Aviation Administration (FAA), domestic passenger air travel growth is expected to average 2.0% over the next 20 years [1]. Aviation demand is driven by economic activity, where the growing U.S. and world economies provide the foundation for long term aviation growth, making efficient and safe air transportation operations more important than ever before. With such a large increase in the volume and complexity of anticipated operations, maintaining or improving safety for all aviation operations is of paramount importance. To meet this objective, global efforts have been underway to modernize aviation systems to address current and future air transportation challenges.

[1]Graduate Research Assistant, Aerospace Systems Design Laboratory, Georgia Institute of Technology, Atlanta, GA
[2]Research Engineer II, Aerospace Systems Design Laboratory, Georgia Institute of Technology, Atlanta, GA
[3]Senior Research Engineer, Aerospace Systems Design Laboratory, Georgia Institute of Technology, Atlanta, GA
[4]S.P. Langley NIA Distinguished Regents Professor, Director of Aerospace Systems Design Laboratory, Georgia Institute of Technology, Atlanta, GA

These global modernization efforts include the FAA's Next Generation Air Transportation System (NextGen) [2] portfolio in the U.S. and the Single European Sky Air traffic management (ATM) Research (SESAR) [3] program in Europe. All global modernization efforts are long-term plans motivated by increasing efficiency and capacity of airspace systems, while maintaining or improving safety.

Historically, accidents or incidents have been the primary triggers for detecting problems and developing mitigation strategies. In the past 25 years, the aviation industry has moved toward proactive and predictive approaches to safety. A proactive approach to safety involves identifying potential unsafe events before they manifest as accidents or incidents such that mitigation strategies may be developed to prevent the occurrence of accidents or incidents related to the unsafe events [4]. To successfully implement a proactive approach to safety, records must exist of the current states of operations. Taking safety analysis a step further, a predictive approach to safety involves monitoring data obtained from routine operations in addition to accidents and incident data and reports to detect potential negative future outcomes. Thus, reactive approaches to safety focuses on prevention of accident or incident *recurrence*, while proactive and predictive approaches to safety focus on prevention of accident or incident *occurrence*. This shift toward proactive and predictive approaches is key to maintaining and improving aviation safety in the future as aviation systems modernize, become more complex, and experience an increase in volume of operations. To maintain or improve safety, recognition and timely mitigation of previously unknown safety risks and threats are necessary. In light of the need to enable proactive and predictive approaches to safety, NASA's Aeronautics Research Mission Directorate's (ARMD) has identified In-time (Real-time) System-wide Safety Assurance (ISSA) as one of its Strategic Thrusts [5–7]. The vision of this Strategic Thrust is "the ability to predict, detect, and mitigate emergent safety risks throughout aviation systems and operations", which requires the development of new classes of offline methods to detect emergent, or previously unknown, safety risks [5]. In this context, *offline* refers to analysis of stored historical data as opposed to online analysis of streamed data in real time or near-real time [6].

As a product of aviation modernization efforts, new data-generating technologies have been introduced in aviation systems, which has resulted in more sources, volume, and availability of aviation operational data than ever before [2]. Most of the aviation operational data is collected with a focus on operations in and around aircraft as they fly through the airspace system [8]. Automatic Dependent Surveillance-Broadcast (ADS-B) technology has been deployed to enhance aviation safety and efficiency by enabling aircraft to determine their position with respect to other similarly-equipped aircraft, using satellite, inertial, and radio navigation [9]. ADS-B Out periodically emits (at approximately 1 Hz) the aircraft's position, along with other relevant parameters, to ground stations and other equipped aircraft [9]. The expansion of ADS-B technology has enabled open-source flight tracking data to become publicly available, which makes analyzing aircraft operations at a larger scale more feasible [10]. This explosion in the volume of aviation operational data recorded may be exploited to provide many new opportunities related to safety analysis [11–14]. In particular, the vast amount of aviation operational data may be exploited to discover previously unknown safety risks in existing aviation systems arising as a result of the implementation of new concepts of operation related to modernization efforts [15].

Advances in modern machine learning techniques have significantly contributed to the use of data-driven techniques to gain insight from operational data and improve aviation safety. The discovery of previously unknown safety risks in aviation data falls under the scope of knowledge discovery and information extraction. Specifically, data mining methods play a significant role in understanding the National Airspace System (NAS), specifically contributing to the development of capabilities that both mitigate current unsafe operations and anticipate future unsafe operations [14]. Anomaly detection is an important step in improving safety of the air transportation system as evidenced by the recent literature on this topic [4, 16–19]. Anomaly detection utilizing quantitative time-series data and unsupervised or semi-supervised machine learning techniques has been particularly popular [17, 19, 20]. The utilization of machine learning techniques to detect anomalies presents many benefits over methods that rely on detecting exceedances in a specified set of parameters. Exceedance detection, one of the most common methods of analyzing flight operational data [14, 20, 21], is concerned with the deviation of a metric or multiple metrics beyond an established threshold within a specified time interval [20]. This method has a number of limitations, one of which being its reliance on a predefined criteria to detect anomalies, which leaves emergent, or previously unknown, safety risks undetected [22]. Machine learning techniques help alleviate this issue, leading to their application for the detection of anomalies in aviation operational data becoming more popular.

Operations within the terminal airspace, in particular, are known to greatly impact both individual flight-level and entire airspace system-level safety and efficiency [23]. Hence, 80% of all accidents and incidents are known to occur within the terminal airspace system [24]. This can be explained by the complex nature of the terminal airspace system itself, which experiences high traffic density of converging and diverging aircraft, interdependent utilization of both airport and airspace resources, and is generally a more constrained airspace featuring complex traffic dynamics [10]. These features make the task of trajectory anomaly detection within the terminal airspace challenging. In addition, in the 10-year period from 2004 through 2013, 47% of fatal accidents and 40% of fatalities have occurred during the final approach or landing phases of flight [25]. Consequently, developing data-driven methods for trajectory anomaly detection on arriving aircraft operational data is of primary interest and the main focus of this paper.

The remainder of the paper is organized as follows: Section 2 provides some relevant background on data-driven methods for aviation anomaly detection, specifies the ADS-B data as the data source to be utilized in this work, and introduces the overarching research objective. Section 3 reviews the existing literature on anomaly detection methods with a focus on two types of anomalies, spatial anomalies and energy anomalies, as introduced in Section 2. Section 4 presents the technical framework developed to identify spatial and energy anomalies utilizing ADS-B trajectory data. It also presents an approach to exploring the relationship, if any exists, between spatial and energy anomalies. Section 5 presents the implementation and results of the proposed framework and provides a discussion on the insights gained. Finally, Section 6 concludes the paper and offers avenues for future work.

2. Background and Motivation

Performing anomaly detection relies on the selection and specification of relevant data and metrics. The literature related to aviation anomaly detection may be split into two categories based on the data source utilized: Flight Operational Quality Assurance (FOQA) data and ADS-B data. FOQA data is collected as part of a structured routine data collection and analysis program, where Flight Data Recorders (FDR) on-board aircraft record massive amount of data for each flight. The amount of operational data recorded includes anywhere between 80 to 2,000 metrics at a sampling rate of 0.25 to 8 Hz [26]. It is noted that FOQA data typically contains ADS-B data metrics in its plethora of metrics, yet a distinction is made between the two due to the difference in their availability for research purposes. ADS-B technology provides messages containing satellite-derived position and speed metrics such as latitude, longitude, altitude, heading, horizontal speed and vertical speed. The flight trajectory data available in ADS-B messages is the core information that is utilized by air traffic management (ATM) systems as a basis for distributing flight information to relevant airlines and air traffic control (ATC) units, facilitating timely coordination between sectors and units, correlating flight data with tracks, monitoring the adherence of an aircraft with its assigned route, and detecting and resolving conflicts. A prominent issue aircraft face during approach and landing is managing energy [27]. An aircraft may possess excess energy due to it being too high on glide slope (high potential energy) or in instances where a high tail wind is experienced (high kinetic energy)[27]. Energy state awareness and energy management are critical components of a safe approach and landing [28]. Due to the open-source accessibility of ADS-B data and the ability to derive metrics related to the position and energy management of an aircraft, many efforts have focused on exploiting ADS-B data to detect anomalies.

As discussed, data-driven methods to detect anomalies in flight trajectory data have become increasingly important in recent years as air traffic volume has continued to rapidly increase. Human operators and analysts are simply not able to properly assess trajectory data for emergent safety risks due to the high volume and dimensionality of the data they are presented with. Therefore, it is of interest to develop anomaly detection methods to narrow the scope of trajectories to be further analyzed by human experts. Further, data-driven anomaly detection methods are better equipped to provide operational insights related to interdependencies in high dimensional time series data. Overall, the objective of anomaly detection methods is often to support aviation operators and decision-makers in facilitating a safe and efficient airspace.

Due to the importance of aircraft position and energy management, *spatial metrics* and *energy metrics*, respectively, are most often the metrics considered when performing anomaly detection utilizing ADS-B data. Typically, an anomaly detection framework will focus on detecting anomalies utilizing either spatial metrics *or* energy metrics. Spatial metrics are those describing an aircraft's position such as latitude, longitude, altitude, or heading. Energy metrics are those derived from the position and speed data that indicate the energy state of the aircraft, such as specific potential energy, specific kinetic energy, specific total energy, and their respective rates. Detecting anomalies utilizing the spatial metrics of an aircraft typically occurs through *air traffic flow* identification. While the precise definition of an air traffic flow depends on the application, it is generally considered a pattern of air traffic in the spatial and, in some applications, temporal dimensions. A large portion of aviation research focuses on the spatial dimension of air traffic flows [11, 29–38]. However, it is occasionally observed that there are trajectories of flights that do not appear to belong to any specific air traffic flow.Much of the literature related to identifying air traffic flows reports a certain percentage of the trajectories detected as outliers, where these trajectories may be considered *spatial anomalies*.

For the purpose of this work, the concept of a *spatial anomaly* is defined as *a trajectory whose spatial metrics do not conform to an identified set of air traffic flows representing standard spatial operations*. On the other hand, the concept of an *energy anomaly* is defined as *a trajectory within an air traffic flow whose energy metrics do not conform to standard energy operations*. Figure 1 shows a notional representation of spatial and energy anomalies per the definitions provided above. The spatial anomaly does not conform to any of the identified air traffic flows (indicated by the different colors) and the energy anomaly exhibits unusual energy patterns compared to the other flights within the flow.

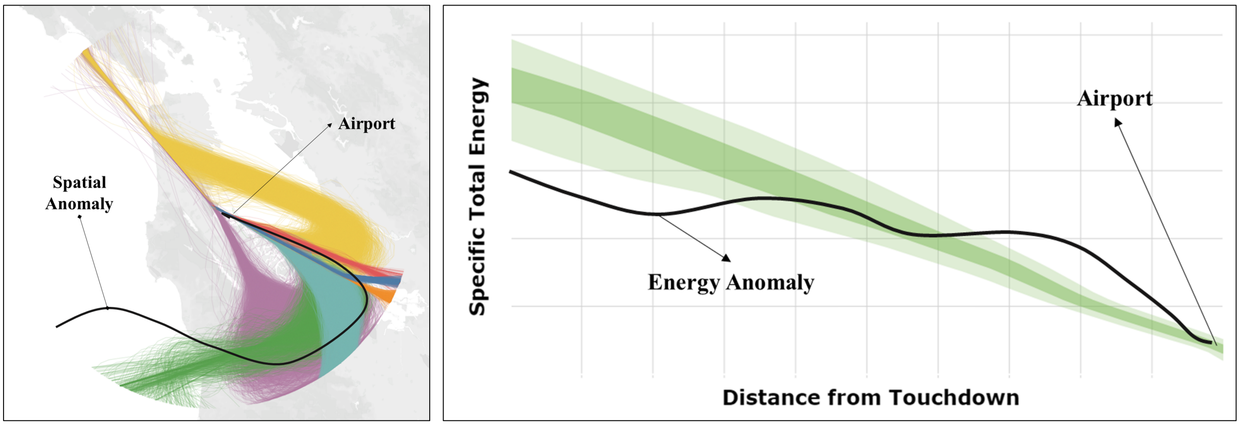

Figure 1: Notional depiction of spatial and energy anomalies among flight trajectory data

A distinction is made between spatial and energy anomalies due to the relative scarcity of approaches that consider both spatial and energy metrics in the same analysis and the differing dimensions the anomalies are discovered in. In the context of detecting *energy anomalies*, air traffic flow identification is sometimes utilized as a data preprocessing step to identify groups of trajectories following a similar spatial path. Trajectories belonging to the same air traffic flow tend to correspond to a specified structured operation defined by ATC within the terminal airspace, i.e. specified headings, altitudes, speeds, and/or waypoints. For each structured operation, the approach and climb paths mandated by ATC typically have distinct energy profiles. Thus, unless an energy anomaly detection method utilizes the additional spatial context related to the air traffic flow to which the trajectory belongs, the method may yield more conservative results. For instance, as noted by Deshmukh, applying energy anomaly detection methods to an entire terminal airspace data set without contextual spatial information may result in anomalous flights remaining undetected as their states are embedded within the entire data set [32].

From the literature, it is unclear whether trajectories that are spatial anomalies, when discovered during air traffic flow identification, are of any further consequence with respect to their energy management. Similarly, there is also no consensus on the spatial behavior of flights with poor energy management. It is asserted that a unified framework to detect both spatial and energy anomalies as part of the same analysis is required in order to understand the relationships and interdependence between these two types of anomalies. This relationship is of interest to both ATC and pilot or air crews for an accurate assessment of risk and mitigation of potentially unsafe operations. Considering the above observations leads to the overarching research objective of this work, which is stated as follows:

Demonstrate a novel, unified framework for the detection of spatial and energy anomalies in ADS-B trajectory data (and associated derived metrics) and explore the relationship between the two types of anomalies in the terminal airspace to better assess the potential for unsafe operations.

3. Literature Review

As presented in Section 2, a distinction is made between spatial anomalies and energy anomalies based upon the types of metrics utilized in the analysis. A review of the literature related to anomaly detection methods that may be classified as spatial anomaly detection methods and energy anomaly detection methods is presented.

3.1. Spatial Anomaly Detection

Often, spatial anomaly detection is a byproduct of an air traffic flow identification task. Spatial air traffic flow identification is often performed to support a more in-depth analysis of the efficiency of airspace operations. Specifically, spatial air traffic flow identification techniques have been developed to support performance assessments, airspace monitoring efforts, more effective air traffic flow management, and system-level ATM decision-making [23]. In addition, as mentioned in Section 2 air traffic flow identification may be implemented as a pre-processing step to generate data sets for which to apply energy anomaly detection methods.

Air traffic flows are typically identified by leveraging trajectory clustering techniques. *Clustering* may be defined as "the task of dividing the population or data points into a number of groups such that data points in the same groups are more similar" [39]. The *trajectory clustering* task is generally formulated as an unsupervised machine learning problem. The task involves partitioning a collection of trajectories into similar groups to define air traffic flows [40]. Several different trajectory clustering methods have been employed to identify air traffic flows in previous works. Eckstein [41] presented an automated flight track taxonomy to identify air traffic flows within a terminal airspace by combining Principal Components Analysis (PCA) and k-means clustering. A limitation of the k-means clustering algorithm is that it does not identify any outliers (detect anomalies) during the clustering. Rehm [42] utilized a single linkage hierarchical clustering technique to identify air traffic flows, where a grid-based approach to outlier (anomaly) detection is utilized prior to the clustering. Enriquez & Kurcz [43] developed an air traffic flow detection algorithm based on spectral clustering. Enriquez [40] built on the work of Enriquez & Kurcz [43], and presented a method to identify temporally persistent flows in the terminal airspace utilizing spectral clustering methods, where irregular air traffic is identified such that corresponding trajectories are detected as being anomalous. Gariel et al. [29] proposed two methods for identifying air traffic flows: (1) way-point-based trajectory clustering and (2) trajectory-based clustering via PCA. Both of Gariel et al.'s proposed methods leveraged the DBSCAN (Density-Based Spatial Clustering of Applications with Noise) algorithm, which is robust to noise in that outliers (anomalies) that do not belong to an identified air traffic flow are officially designated as outliers by the algorithm [44]. Conde Rocha Murca et al. [23] presented a data mining framework to characterize air traffic flows in metroplex systems, where DBSCAN is utilized to learn the typical network of nominal arrival and departure flows. Olive & Morio [9] proposed a method to identify air traffic flows specifically in the terminal airspace, which recursively applies DBSCAN to cluster significant trajectory points, and then builds a dependency tree that is ultimately utilized to label trajectories as belonging to an air traffic flow or as outliers (anomalies). Additionally, Olive & Basora [45] demonstrated the use of a Gaussian Mixture Model for trajectory clustering. In other work, Olive & Barsora [35] utilized the Hierarchical DBSCAN (HDBSCAN) algorithm to identify air traffic flows and anomalies for en-route trajectories. Basora et al. [46] presented a trajectory clustering framework utilizing HDBSCAN to analyze air traffic flows. Additionally, Tanner & Strohmeier [38] applied HDBSCAN to cluster surface trajectories for runway utilization assessments. Most recently, Olive et al. [36] demonstrated the utilization of deep clustering techniques to identify and characterize air traffic flows. Finally, as mentioned, air traffic flow identification (through trajectory clustering) is often utilized as a data pre-processing step to energy anomaly detection to prevent overly conservative energy anomaly detection models from being developed [32]. For instance, Puranik [4], Deshmukh [32], and Olive & Basora [37] utilized DBSCAN to identify air traffic flows and outliers prior to performing energy anomaly detection. However, the context in which the spatial outliers are considered in the energy anomaly detection analysis is unclear, if they are considered at all.

To perform the methods of trajectory clustering presented, similar data instances are grouped into clusters according to a defined *distance function*. It has been noted that one of the most critical components of trajectory clustering is the definition of an appropriate distance function [9]. Trajectories are functional in nature, not independent data points, so the selection of an appropriate distance function is often less straightforward [36]. The Euclidean distance is most often utilized to compute the distance between two n-dimensional trajectories [37]. Provided two n-dimensional trajectory vectors, T^i and T^j containing x and y features, where $T^i = [(x^i_1, y^i_1), (x^i_2, y^i_2), ..., (x^i_n, y^i_n)]$ and

$T^j = [(x_1^j, y_1^j), (x_2^j, y_2^j), ..., (x_n^j, y_n^j)]$ the Euclidean distance computation proceeds as:

$$D_{ED}^{i,j} = \sqrt{\sum_{k=1}^{n} [(x_k^i - x_k^j)^2 + (y_k^i - y_k^j)^2]},$$

where x_k and y_k indicate an aircraft's horizontal position related to longitude and latitude measurements, respectively, at point k in the sequence of points making up the aircraft's trajectory.

Due to the convergence/divergence of trajectories within the terminal airspace, the distances computed between trajectory points closest to the airport are relatively small, regardless of the air traffic flow to which they belong. This is due primarily to runway configuration and assigned arrival and departure runways. On the other hand, distances between trajectories within the same flow may be relatively large between trajectory points considered at the terminal airspace's defined border. The distance between these trajectory points depends on the density of the air traffic flows and the radius around the airport. The uneven distribution of distances as aircraft arrive at or depart from the airport may skew the classification, resulting in misclassification of trajectories or inadequate air traffic flow identification. Therefore, it is noted that the utilization of the Euclidean distance to cluster trajectory records may not be adequate [9, 47].

Other distance functions have been proposed to address the issues related to the limitations of the Euclidean distance, such as the Symmetrized Segment Path Distance (SSPD) [47]. Basora et al. [46] compared the utilization of the Euclidean distance versus the SSPD to cluster en-route trajectories and found that the utilization of the SSPD results in more precise identification of air traffic flows. Yet, there exists a significant trade-off in terms of computation time, with a large computation time likely to prevent subsequent applications of the algorithm in real-time. Gariel et al. [29] and Olive & Morio [9] presented iterative point-based clustering methods to identify air traffic flows. However, clustering of the entire trajectory record or multi-point segments at once is often a requirement to enable the extension of clustering methods to real-time applications. While real-time application is not the focus of this work, aviation organizations have set goals to ultimately deploy novel anomaly detection methods in real-time or near-real-time [6]. Corrado et al. [48] proposed the utilization of a weighted Euclidean distance function to cluster trajectories and demonstrated that certain weighting schemes resulted in a more accurate clustering. The weighted Euclidean distance computation proceeds as:

$$D_{WED}^{i,j} = \sqrt{\sum_{k=1}^{n} w_k[(x_k^i - x_k^j)^2 + (y_k^i - y_k^j)^2]}.$$

The weighted Euclidean distance introduces an additional trajectory point weighting term, w_k, which enables an "importance" to be assigned to each trajectory point. The importance may be determined according to the position in the sequence of points relative to the airport's location or bounds of the airspace.

Distance functions such as the Euclidean distance and weighted Euclidean distance require re-sampling all trajectories to standardize the length of the trajectories such that each trajectory is represented by an n-dimensional vector of points [37]. In the context of spatial anomaly detection within the terminal airspace, the data set of interest contains trajectories of flights within some radius (or bounding box) of the airport [9, 40–43]. Because these trajectory records are typically not of the same length, re-sampling is necessary. Both distance-based and time-based re-sampling methods have been applied to generate an n-dimensional vectors of points to represent a trajectory operating within some radius of the airport.

Some studies have also focused on identifying and predicting go-arounds [49–51]. Go-arounds are a well-practiced, yet relatively rare procedure, often undertaken due to approach stability or air traffic control considerations [49]. Due to the their rarity, they may be considered a specific category of spatial anomalies as they do not belong in the category of nominal operations within the terminal airspace. Go-arounds are generally easily detected in historical data utilizing logic-based assessments of metrics such as vertical rate, altitude, and cumulative ground track distance from touchdown [49].

3.2. *Energy Anomaly Detection*

Recently, Basora & Olive [18] published an up-to-date review on the recent advances in anomaly detection methods applied to aviation data, in which several of the methods reviewed may be classified as energy anomaly detection methods. Generally, when "anomaly detection" is referred to in aviation literature, the implication is that the anomaly detection methods are applied to either energy metrics derived from trajectory data, or those utilizing FOQA data. While anomaly detection utilizing FOQA data is not the focus of this work, it is important to introduce a few relevant and prominent methods as ADS-B trajectory data energy anomaly detection methods take inspiration from these methods. Methods utilizing ADS-B trajectory data for anomaly detection in energy metrics are reviewed. One of the first and most effective anomaly detection methods is multiple kernel anomaly detection (MKAD), developed by Das et al. [16], which detects anomalies in heterogeneous sequences of both continuous and discrete FOQA data features. The novelty of the MKAD method is the ability to successfully handle the heterogeneous sequences of both continuous and discrete FOQA data features [16], where the ADS-B trajectory data energy metrics are simply continuous sequences of metrics indicating the energy states of the aircraft over time. Li et al. [52] presented ClusterAD, a clustering-based anomaly detection method leveraging the DBSCAN algorithm to detect anomalies in FOQA data. ClusterAD relies on an a-priori specification of the desired fraction of anomalies to be detected, which is controlled by certain DBSCAN inputs parameters [52]. Li et al. [17] extended ClusterAD to develop a ClusterAD-Flight, which identifies anomalies at the flight level. Additionally, Li et al. [21] presented ClusterAD-DataSample, which leverages Gaussian Mixture Model (GMM) clustering to detect instantaneous anomalies in FOQA data. Sheridan et al. [22] presented a DBSCAN-based method for anomaly detection in the approach phase.

Puranik et al. [28] defined energy metrics to be computed from trajectory data records. Puranik et al. [53], Puranik & Mavris [19, 20, 54], and Puranik [4], utilized these energy metrics to detect flight-level and instantaneous anomalies in general aviation operations. Specifically, Puranik & Mavris [19] and Puranik [4] presented a method leveraging the DBSCAN algorithm and Support Vector Machines (SVMs) to detect anomalies in departing and arriving general aviation flights. In the context of commercial operations, Kim & Hwang [33], Deshmukh & Hwang [11, 30], and Deshmukh [32] proposed TempAD, an algorithm designed to provide formulas related to the bounds of normality that are easily interpreted in natural languages, where this method utilizes DBSCAN to identify air traffic flows as a data pre-processing step. TempAD has been utilized to detect anomalies in the vertical dimension (altitude), the speed dimension (ground speed), and energy metrics such as specific total energy and specific potential energy rate [11, 30, 32]. Jarry et al. [55] presented a method to detect trajectories as being anomalous in energy metrics utilizing a Functional Principal Components Analysis (FPCA)-based approach combined with a sliding window and outlier scoring. Jarry et al. [27] expanded the previous work by leveraging FPCA combined with Hierarchical DBSCAN (HDBSCAN) and Global-Local Outlier Score from Hierarchies (GLOSH)to detect anomalies in energy metrics within the terminal airspace. One may note that the utilization of DBSCAN-based methods is very prevalent within the aviation anomaly detection literature.

Similar to spatial anomaly detection, the anomaly detection methods reviewed related to energy anomaly detection also rely on re-sampling the varied-length trajectories to be n-dimensional vectors of points. However, unlike spatial anomaly detection, the scope of the data utilized often depends on a distance- or time-based threshold for cutting off the data. For instance, Das et al. [16] utilized the FOQA data recorded below 10,000 ft mean sea level and Li et al. [52] utilized the FOQA data recorded during the final six nautical miles of the approach. Additionally, Deshmukh & Hwang considered trajectory data 20 nautical miles from touchdown [11] and Jarry et al. [55] considered trajectory data 25 nautical miles from touchdown, both utilizing distance-based thresholds to scope the data set.

3.3. Limitations of Existing Literature

As introduced in Section 2, a distinction should be made between spatial and energy anomalies. In the context of energy management, it is unclear whether the operations of trajectories that do not conform to standard spatial operation should receive any special consideration. Conversely, the impact of poor energy management on spatial operations is unknown. Knowledge of the relationship and/or interdependencies between spatial and energy anomalies may enable ATC and other aviation operators to make ATM decisions that are more robust to safety risks. Within the aviation literature, a unified framework does not exist to perform both spatial and energy anomaly detection. Historically, these two types of anomaly detection methods are considered separately. Spatial anomaly detection is generally only considered in the context of air traffic flow identification, where outliers (anomalies) are a byproduct of applying many of the trajectory clustering techniques. Energy anomaly detection methods do not appear to make any special consideration of spatial anomalies, if considered at all. The implication of spatial anomalies when applying energy anomaly detection methods has not yet been determined, though is of interest to aviation system operators in the context of improving safety.

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 17 February 2021 doi:10.20944/preprints202102.0396.v1

4. Methodology

The purpose of the proposed methodology is to provide a novel, unified framework for the detection of spatial and energy anomalies in ADS-B trajectory data (and associated derived metrics) and to provide a structure for which to explore the relationship between the spatial and energy anomalies detected utilizing this framework. The framework consists of three main parts – a data processing procedure, a spatial anomaly detection method, and an energy anomaly detection method. While the spatial and energy anomaly detection methods used within the framework are clustering-based methods, there is no restriction on the type of anomaly detection method that may be utilized for either spatial or energy anomaly detection.

Data processing occurs such that spatial and energy metrics are appropriately sampled and derived for the spatial and energy anomaly detection tasks, respectively. As mentioned in Sections 2 and 3, there typically exists differences in energy profiles between differing structured operations in the terminal airspace, generally corresponding to the different identified air traffic flows. Therefore, spatial anomaly detection is performed prior to energy anomaly detection. The steps involved in the development of this novel framework for detection of spatial and energy anomalies in ADS-B trajectory data are displayed in 2, with each of the primary steps, namely (i) data processing, (ii) spatial anomaly detection, and (iii) energy anomaly detection, being discussed in detail below. Finally, a process for exploring the relationship, if any, between spatial and energy anomaly characteristics is also discussed later in this section.

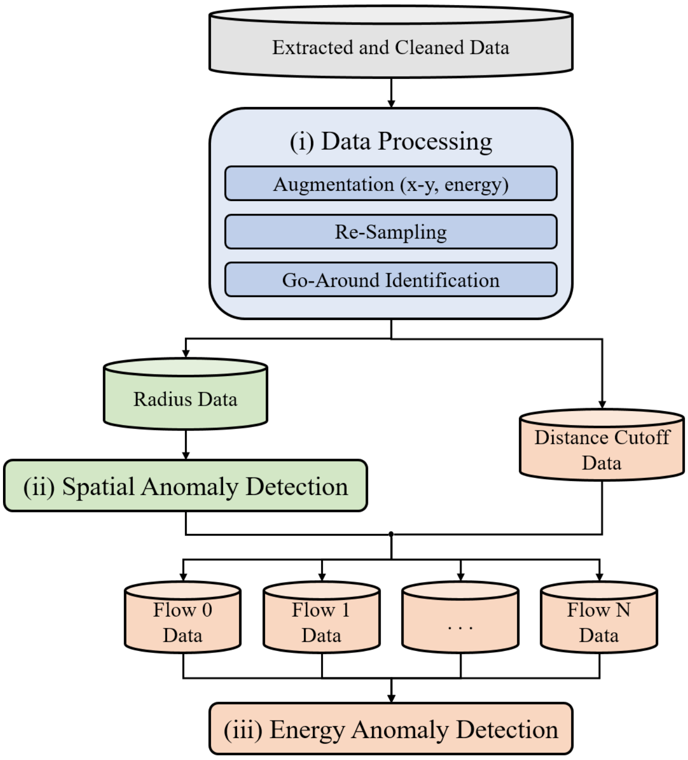

Figure 2: Framework to detect spatial and energy anomalies in ADS-B trajectory data

4.1. Data Processing

Provided a set of ADS-B trajectory data for flights operating within a terminal airspace during a specified time period, additional metrics must be computed to enable both spatial and energy anomaly detection. Spatial anomaly detection requires that each longitude and latitude pair be projected onto a Universal Transverse Mercator (UTM) coordinate system to generate Cartesian coordinates for spatial clustering. Energy anomaly detection requires the computation of the appropriate energy metrics. This paper focuses on three primary energy metrics in which to detect anomalies: specific potential energy (SPE), specific kinetic energy (SKE), and specific total energy rate (STER).

The specific potential energy is taken as the height above ground level of the aircraft. The specific kinetic energy is computed as:

$$SKE = \frac{V^2}{2g},$$

where V is the ground speed of the aircraft available in the ADS-B data and g is the gravity constant. The specific total energy rate requires the computation of the specific total energy (STE), which is the sum of the specific potential energy and specific kinetic energy:

$$STE = SPE + SKE.$$

The specific total energy rate is then computed as:

$$STER = \frac{STE_{i+1} - STE_i}{\Delta t},$$

where Δt is the time difference between two consecutive records, $i + 1$ and i.

As noted in Section 3, the scope of the ADS-B trajectory data utilized for spatial and energy anomaly detection often differs. Spatial anomaly detection is typically performed considering all trajectories within some radius (or bounding latitude/longitude) of the airport, i.e. the record of a single trajectory begins once the aircraft enters the specified radius from the airport and concludes upon touchdown. All trajectories within the data set utilized for spatial anomaly detection will have different maximum cumulative ground track distance from their respective touchdown points and different temporal lengths. On the other hand, energy anomaly detection is typically performed considering a data set that has been cut off at some distance- or time-based threshold, i.e. 20 nautical miles cumulative ground track distance from the touchdown point or the final 10 minutes of flight. In this work, a distance-based cutoff is proposed to mitigate the effects of potentially differing approach speeds for different aircraft type.

As mentioned in Section 3, spatial anomaly detection methods typically operate on trajectory records of a standardized n-dimensional length, and the same is true of the energy anomaly detection methods reviewed. Therefore, a re-sampling of the data to form n-dimensional trajectory records is often a necessary processing step before performing spatial or energy anomaly detection. There are two primary re-sampling methods: distance-based and time-based. In the case where aircraft maintain a relatively constant velocity, distance-based and time-based re-sampling methods generate a nearly identical set of points. However, within the terminal airspace, aircraft velocity is rarely constant for extended periods of time, particularly during the approach phase where it gradually decreases. A time-based re-sampling would result in an uneven balance of points closer to the airport. Therefore, a distance-based re-sampling method is selected. A uniform re-sampling occurs based on the cumulative ground track distance of an aircraft from its touchdown point. This results in the re-sampled points being spatially more evenly distributed. After re-sampling of the two data sets, the selected features are stacked to create feature vectors for utilization with the clustering algorithms. As such, the data processing step produces two re-sampled data sets to be utilized for anomaly detection: a data set containing feature vectors corresponding to trajectories operating within a specified radius of the airport and a data set containing records cut off at a specified cumulative ground track distance from touchdown. Both data sets will be of the same size and contain trajectories of the same flights, yet the data set containing records cut off at a specified cumulative ground track distance from touchdown can be thought of as having a higher "resolution".

An additional data processing step includes the identification go-arounds. The concept of go-arounds as a category of spatial anomalies is introduced in Section 3. Go-arounds are almost always either excluded from analysis or detected by a spatial anomaly detection algorithm as being anomalous. In this framework, go-arounds are identified as trajectories where, when an aircraft's cumulative ground track distance from touchdown is greater than 25 nautical miles, its height above ground level reaches below 2,500 ft and the vertical rate reaches above 5 ft/min. If left within the spatial anomaly detection data set, go-arounds should always be detected as anomalous by a robust spatial anomaly detection algorithm, as it is known that go-arounds are not associated with standard terminal airspace operations. Because go-arounds are easily identifiable within a trajectory data set and are known spatial anomalies, it is most appropriate to remove the go-arounds from the data set such that only trajectories remain for which no spatial anomaly label is specified a-priori. This allows the algorithm to search for and find hitherto unknown anomalies. Due to the existence of an easily obtainable spatial label for go-arounds, the assessment of go-arounds in the context of anomaly detection and safety should be performed separately. After the removal of go-arounds from the data set, feature vectors containing relevant spatial and energy metrics for the spatial and energy anomaly detection tasks, respectively, are generated.

4.2. Spatial Anomaly Detection

Spatial anomaly detection typically occurs as a byproduct of air traffic flow identification. Hence, in this work, an air traffic flow identification algorithm is utilized to perform spatial anomaly detection. The metrics utilized for spatial anomaly detection are the Cartesian coordinates resulting from the projection of longitude and latitude onto a UTM coordinate system. Due to the increased prevalence and capability of the HDBSCAN [56] algorithm within the aviation literature [27, 46, 57], HDBSCAN is selected as the clustering algorithm to perform the spatial anomaly detection task. The HDBSCAN algorithm extends the widely-utilized DBSCAN algorithm by converting it to a hierarchical clustering algorithm, where it ultimately extracts a flat clustering (set of clusters without any explicit structure that relates the clusters to each other) based on the stability of the identified hierarchical clusters [58]. The algorithm's primary advantage over DBSCAN is in identifying clusters of varying density, which is relevant for applications in the terminal airspace where flows often have varying densities, i.e. some flows appear very tight, while others are more spread out. HDBSCAN requires a single input parameter, $minSamples$, that sets the minimum cluster size to perform clustering [56]. Additionally, an optional smoothing parameter, $minClusterSize$ may be set that provides a measure of how "conservative" the clustering should be [58]. The larger the smoothing parameter value, the more points declared as outliers and the clusters restricted to more dense areas [58]. The readers are referred to Campello et al. [56] for further details on the algorithm.

Clustering algorithms rely on the definition of a distance function. In previous work by the authors [48]), the effectiveness of a weighted Euclidean distance function to identify air traffic flows is demonstrated. The results indicated a more robust identification of outliers. Therefore, the recommended weighted Euclidean distance function is selected to be utilized with HDBSCAN to assign each trajectory to an air traffic flow or detect the trajectory as being anomalous. If a trajectory is detected as anomalous, the closest air traffic flow is also noted, i.e. the air traffic flow whose centroid is the shortest distance (computed utilizing the selected weighted Euclidean distance) from the anomalous trajectory. Thus, the final product of spatial anomaly detection consists of two labels for each trajectory: one indicating whether the trajectory is spatially nominal or spatially anomalous, and another indicating the air traffic flow to which the trajectory most closely belongs.

4.3. Energy Anomaly Detection

Performing energy anomaly detection on the entire terminal airspace data set without contextual spatial information, such as the air traffic flow to which trajectories belong, may result in anomalous flights remaining undetected as their states are embedded within the entire data set. In other words, while a trajectory may appear nominal with respect to the entire set of trajectories, it may actually be considered anomalous with respect to other trajectories following a similar spatial path, or corresponding to the same structured operation within an airspace [32]. Therefore, for each of the air traffic flows identified, an energy anomaly detection method is applied separately. However, it is not specified within the existing literature if there is additional consideration of a trajectory's energy profile after it has been detected as a spatial anomaly in the air traffic flow identification step. Proper energy management is paramount to the safety of approaches and landings, so it is important to assess the energy states of all trajectories, even those that have already been detected as being spatially anomalous. Further, if spatially anomalous trajectories' energy profiles are considered when detecting energy anomalies, it is not clear how they are considered, which is significant in the context of evaluating the energy anomalies that are ultimately detected. It is of interest to ATC and other aviation operators, for instance, to know whether a trajectory that does not conform to standard spatial operations has a higher likelihood of experiencing off-nominal energy states, as this may pose a safety risk. While an exploration of the relationship between spatial and energy anomalies offers opportunities to obtain actionable insights as it relates to ATM decision-making, such effort is currently lacking in the literature. Overall, it is important to understand the context in which spatial anomalies are placed when proceeding with energy anomaly detection such that aviation operators may be aware of potential impacts or risks of certain instructions that result in an aircraft deviating from standard spatial operations. From the literature it is unclear whether a trajectory that has been detected as a spatial anomaly is considered when detecting energy anomalies. However, proper energy management is paramount to the safety of approaches and landings and it is important to assess the energy states of all trajectories, even those that have already been identified as spatially anomalous or nominal. Further, in the context of monitoring efforts to maintain and improve aviation safety, it is important to understand the relationship or interdependencies, if any, between spatial and energy anomalies. To that end, the energy anomaly detection data set is separated into multiple sets based on the closest air traffic flow label. This results in each data set containing some proportion of originally identified spatial anomalies. An energy anomaly detection algorithm is then applied to each data set to detect the energy anomalies.

In the context of clustering energy metrics, it is observed that there is generally one cluster of nominal operations that is identified within each air traffic flow [28]. Due to this property, a clustering algorithm that can identify clusters of differing densities, such as HDBSCAN, is typically not required. Consequently, the widely-utilized DBSCAN algorithm is selected as the clustering algorithm to perform the energy anomaly detection task. The DBSCAN algorithm finds core samples of high density and expands clusters from there, where there is no required a-priori specification of the number of clusters in the data set [44]. Additionally, the algorithm is robust to noise in that outliers (data points that do not meet the criteria to be included in any of the identified clusters) may be excluded from the clustering and assigned an outlier (anomaly) label [44]. DBSCAN relies on two inputs parameters: (i) $minSamples$, which is a minimum number of observations (trajectories) and (ii) ε, which is a distance threshold, where the distance is computed using the specified distance function [44]. While a weighted Euclidean distance is utilized for spatial anomaly detection, the Euclidean distance is utilized for energy anomaly detection. The use of the Euclidean distance is attributed to the relatively equal importance of the energy profile of an aircraft all the way through touchdown, i.e. to detect a high or low energy landing as anomalous, an anomaly detection algorithm would likely require an entire trajectory's energy profile. As indicated, the metrics selected for energy anomaly detection include SPE, SKE, and STER, as combinations of these metrics have been widely used in aviation energy anomaly detection studies [4, 11, 28, 30–32, 53]. To prevent bias in the DBSCAN clustering due to having metrics of varying magnitudes, a normalization is performed prior to applying DBSCAN.

If the $minSamples$ parameter is fixed, there generally exists a monotonic relationship between the ε parameter and the resulting fraction of outliers: as ε is increased, the resulting fraction of outliers is decreased. The utilization of DBSCAN within the aviation literature often requires a set value for $minSamples$ and then varies the ε parameter to reach a pre-defined fraction of outliers [17, 22, 52]. Hence, similar to past studies, to perform the energy anomaly detection task within this framework, DBSCAN is utilized, where the $minSamples$ parameter is fixed and the ε parameter is varied until a pre-defined fraction of outliers is obtained. The energy anomaly detection task is performed for each of the energy anomaly detection data sets introduced previously. The final product the energy anomaly detection exercise is a label for each trajectory indicating whether the trajectory's energy profile is nominal or anomalous.

4.4. Spatial and Energy Anomaly Relationship Exploration

The proposed framework to detect spatial and energy anomalies enables an exploration of the relationship between the two types of anomalies. By comparing the fraction of all spatial anomalies within a flow to the fraction of spatial anomalies identified as energy anomalies within the same data set, it is possible to assess if spatial anomalies are more likely to also be energy anomalies. It is hypothesized that if a trajectory is spatially anomalous, then it is more likely to also be detected as an energy anomaly. In addition, an *energy anomaly score*, or *anomaly score*, is computed for each trajectory within an energy anomaly detection data set to provide more context to the "relative anomalousness" of trajectories operating within the terminal airspace. The anomaly score may be computed as the mean Euclidean distance from a given trajectories' energy metrics to all others in the data set:

$$Score^i = \frac{1}{n}\sum_{j=1, j\neq i}^{n} D_{ED}^{i,j},$$

where $D_{ED}^{i,j}$ is the Euclidean distance between trajectory i's energy metrics and trajectory j's energy metrics.

Four different labels may be assigned to trajectories depending on the category of trajectory they belong to following both spatial and energy anomaly detection. These categories include: (i) nominal (N), (ii) only spatial anomaly (S), (iii) only energy anomaly (E), and (iv) both spatial and energy anomaly (B). The distribution of anomaly scores for trajectories belonging to each category may then be assessed for each flow. To expand on the previous hypothesis, considering only energy-nominal trajectories, it is further hypothesized that spatially anomalous trajectories will, on average, have greater anomaly scores than spatially nominal trajectories. Specifically, it is of interest to determine whether the average anomaly score for only spatial anomalies is greater than the average anomaly score for trajectories that are nominal. If this hypothesis is validated, the implication is that, despite being detected as energy-nominal, only spatial anomalies are "relatively more anomalous" than those that are nominal. This in turn would indicate that a higher degree of caution should still be exercised with respect to trajectories that are veering off a nominal spatial path, even if their energy profiles appear to be nominal.

A one-sided Welch's t-test [59] is conducted to determine whether the average anomaly scores for only spatial anomalies are statistically significantly greater than the average anomaly scores for trajectories that are nominal. Additionally, it is of interest to determine whether the distribution of only energy anomalies' anomaly scores and both spatial and energy anomalies' anomaly scores are generated from the same distribution. Investigating this relationship provides insight into whether a distinction may be made between only energy anomalies and spatial and energy anomalies. It is finally hypothesized that there is a difference between the distribution of anomaly scores of only energy and both spatial and energy anomalies, which would imply that these two categories of anomalies should be assessed independently. A visual inspection is performed to determine whether there is a statistically significant difference between the two different distributions of anomaly scores. The statistical testing of mean anomaly scores and visual inspection of anomaly score distributions help better understand the relationship and potential interdependencies of the four different categories of anomalies.

5. Implementation and Results

The first step in implementing the framework for detection of spatial and energy anomalies in ADS-B trajectory data spatial and energy metrics, respectively, involves extracting and cleaning the data. Therefore, this section begins by detailing the ADS-B trajectory data extraction and cleaning process. Then, the implementation details of each of the steps of the framework are described. Finally, the procedure implemented to explore the relationship between spatial and energy anomalies is discussed.

5.1. ADS-B Trajectory Data Extraction and Cleaning

The complete data extraction and cleaning steps are summarized in Figure 3 and discussed in detail below.

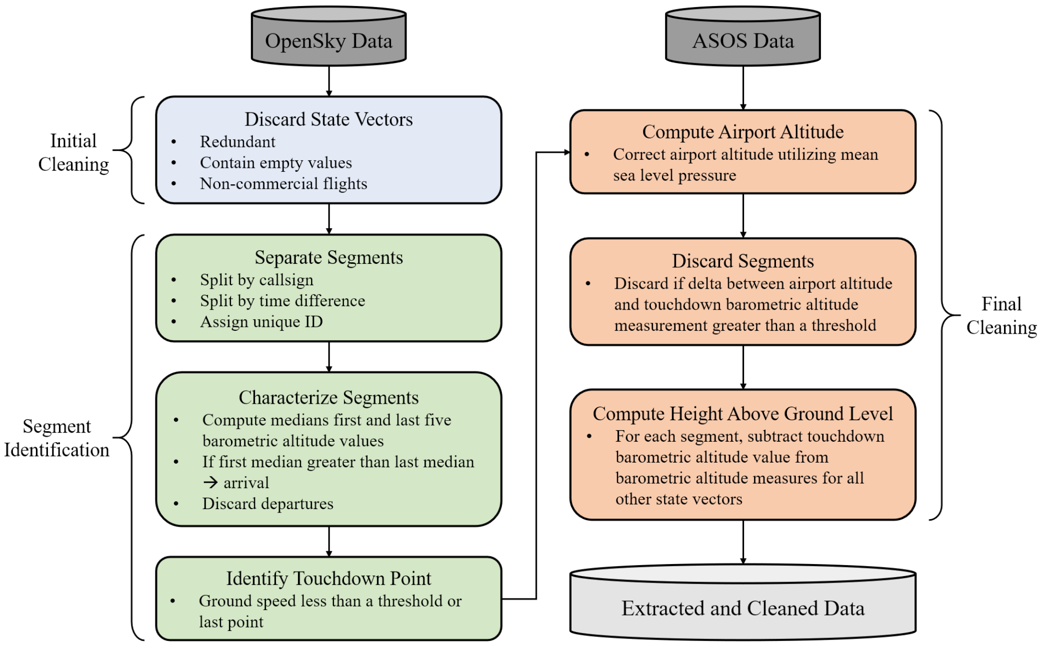

Figure 3: Data extraction and cleaning

The ADS-B trajectory data utilized in this work was extracted from the OpenSky Network's [60] historical database. The OpenSky network is a non-profit association that processes and archives ADS-B data from a global network of sensors [60]. OpenSky data has previously been used by researchers for a diverse range of studies. The Python traffic library [61] is utilized to extract the OpenSky data. Recorded data, or state vectors, available from

 doi:10.20944/preprints202102.0396.v1

OpenSky's historical database, contain timestamps (added on the receiver side, with many receivers equipped with a GPS nanosecond precision clock), transponder unique 24-bit identifiers (icao24), space-filled 8 character callsigns, latitude, longitude, (barometric) altitude (w.r.t. standard atmosphere), GPS altitude, ground speed, true track angle, and vertical speed. The trajectory data available from OpenSky may also be utilized to evaluate the energy metrics mentioned previously using the aircraft's position, velocity, altitude, etc. The OpenSky network has good coverage over Europe and North America and any airport under the coverage can be selected for analysis. For this analysis, trajectories available from the OpenSky network's historical database of flights arriving at San Francisco International Airport (KSFO) between June 1^{st}, 2019 and September 30^{th}, 2019 were considered. OpenSky data was requested via the Python traffic library [61] for a bounding box with a maximum altitude of 25,000 ft and latitude and longitude within 20 nautical miles of the airport location.

It is noted that the OpenSky state vectors provide measurements of barometric altitude, which may vary significantly throughout the day based on current atmospheric temperature and pressure. Therefore, to compute a more accurate height above ground level of the aircraft, atmospheric pressure data for the time period of interest was also extracted. This data was extracted from the Iowa Environmental Mesonet [62], which collects historical Automated Surface Observing System (ASOS) observations at most airports across the United States, including KSFO. The ASOS program is a joint effort of the National Weather Service (NWS), the FAA, and the Department of Defense (DoD). ASOS units are automated sensor suites that are designed to serve meteorological and aviation observing needs and are widely used by meteorologists, climatologists, hydrologists, and aviation weather experts. The ASOS systems serve as the nation's primary surface weather observing network.

To clean the extracted ASOS data required a procedure to fill and interpolate missing values. The OpenSky Network data required a more in-depth procedure. To begin, state vectors that were redundant or contained empty values for latitude, longitude, altitude, heading, or ground speed were discarded. Through an assessment of a state vector's callsign, it was determined whether or not the state vector was associated with a commercial flight. Those not associated with a commercial flight were discarded. At this point, based on time stamps, the ASOS data was fused and appended to the OpenSky state vectors.

Next, state vectors were split by callsign, and then further split into individual trajectory segments where the time difference between two successive state vectors was greater than five minutes. Five minutes was set as the threshold for which to separate flight segments due to the loss of information that would result from five minutes between trajectory measurements within the terminal airspace. Further, it is possible that multiple flights operate with the same callsign in one day, and splitting by a large time difference enables both of these segments to be captured. A unique identifier was assigned to each of the identified trajectory segments. The trajectory segments were then characterized as arrival or departure segments based on the medians of the first and last five altitude measurements. If the first altitude median was greater than the last altitude median, then the trajectory segment was characterized as an arrival. Otherwise, the trajectory segment was characterized as a departure and disregarded for the purpose of this work. The median of the first five altitude measures was taken to be robust to any noise or erroneous measurements that occur.

Then, the touchdown point state vector was identified. The touchdown point was identified as the first state vector where the ground speed reaches below 100 knots, or the last point if no point was recorded to reach below 100 knots. It is noted that this is an approximate method for identifying the touchdown point and can lead to some error, but is observed to be the most robust for this data set. Next, the lowest trajectory segment altitude reached was determined by utilizing the mean sea level pressure from the ASOS data and computing a corrected airport altitude for the touchdown point state vector. The trajectory segment was discarded if the difference between the touchdown altitude and the airport altitude was greater than 100 ft (i.e. the trajectory segment did not reach within 100 ft above ground level). These segments were discarded because enough data does not exist to consider the approach to be complete. The touchdown altitude was then utilized to compute a Height Above Ground Level (HAGL) for each state vector by subtracting the touchdown altitude from all other altitude measurements. Then, the cumulative ground track distance remaining to the touchdown point was computed at each point of the trajectory.

The resulting set of ADS-B trajectory data for flights operating within the KSFO terminal airspace from June 1^{st}, 2019 through September 30^{th}, 2019 contains 63,360 total arrival trajectory segments. The data set contains both domestic and international flights on 129 airlines, where approximately one-third of the flights are operated by United Airlines. Further, the data set contains 90 unique aircraft types, where the most commonly operated is A320 aircraft, closely followed by B737-900 aircraft and B737-800 aircraft. A distribution of the top five most common airlines and aircraft types for the 63,360 flights in this data set is displayed in Figure 4.

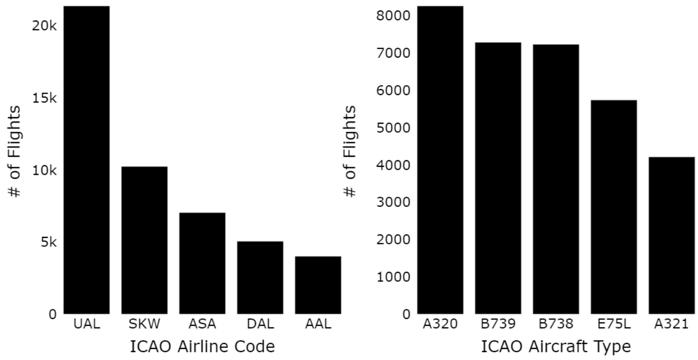

Figure 4: Top five most common airlines and aircraft types

5.2. Data Processing

The data processing task led to the generation of the spatial and energy anomaly detection feature vector data sets to be utilized for the tasks. The extracted and cleaned data set was first augmented with the relevant additional features, such as Cartesian coordinates and selected energy metrics (SPE, SKE, and STER). The spatial anomaly detection data set was generated by re-sampling the extracted and cleaned data to 50 uniformly-spaced points from the maximum cumulative ground track distance from the touchdown point within the 20 nautical mile radius of KSFO to the touchdown point. For this data set, the uniform spacing averages about 0.65 nautical miles ground track distance between each of the 50 re-sampled trajectory points. The energy anomaly detection data set was generated by re-sampling the last 20 nautical miles of flight to 50 uniformly-spaced points. This results in about 0.4 nautical miles ground track distance between each of the 50 re-sampled trajectory points. Within the data set of flights arriving at KSFO between June 1^{st}, 2019 and September 30^{th} 2019, 337 of the 63,360 flights were classified as go-arounds. Identified go-arounds made up only 0.53% of the data set, which indicated a higher go-around rate than 0.29% presented in [63]. As outlined in Section 4, the trajectories of these 337 flights were removed from the data set, leaving 63,023 flight trajectories for analysis.

To complete the data processing task, the data sets for spatial and energy anomaly detection were formatted as feature vectors. For spatial anomaly detection, the Cartesian coordinates were stacked to generate feature vectors for each trajectory. The feature vectors for all trajectories were then aggregated to produce the HDBSCAN clustering input matrix. For energy anomaly detection, the SPE, SKE, and STER metrics were stacked to generate feature vectors for each trajectory. The feature vectors were then separated into energy anomaly detection data sets based upon the closest air traffic flow to which they are assigned. As a result, the number of energy anomaly data sets is equal to the number of identified air traffic flows. The feature vectors for each energy anomaly detection data set are then aggregated to produce DBSCAN clustering input matrices.

5.3. Spatial Anomaly Detection

The *hdbscan* Python library [58] was leveraged to apply the HDBSCAN algorithm to the spatial anomaly detection data set to identify air traffic flows and detect spatial anomalies. As mentioned in Section 4, the HDBSCAN algorithm requires one input parameter, $minSamples$, where an optional smoothing parameter, $minClusterSize$ may be input as well. The *hdbscan* Python library defaults the $minClusterSize$ parameter to the $minSamples$ value if none is inputted [58]. In this study, only the $minSamples$ parameter was explicitly set to a value of 630, which was approximately 1% of all trajectories in the data set. Additionally, a weighting scheme was selected for the utilization of the weighted Euclidean distance with HDBSCAN. Corrado et al. [48] demonstrated the success of their "Weighting 1" weighting scheme on a data set including flights at KSFO between June 1^{st}, 2019 and September 30^{th} 2019. Therefore, this "Weighting 1" scheme was selected. The "Weighting 1" scheme was inspired by a beta probability density function (pdf), and is displayed in Figure 5.

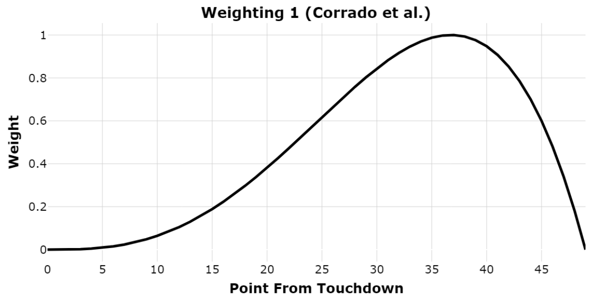

Figure 5: Weighting Scheme Utilized for Spatial Anomaly Detection [48]

After utilizing HDBSCAN with the weighting Euclidean distance to cluster the 63,032 trajectories, 8,082 trajectories were detected as being spatial anomalies (12.8% outliers). Further, seven distinct air traffic flows were identified, as displayed in Figure 6a. Figure 6b presents the proportions of trajectories within the entire data set belonging to one of the seven air traffic flows, having been classified as spatial anomalies by HDBSCAN, or having been identified as go-arounds a-priori. It is noted that the identified air traffic flows were of varying sizes. For instance, flow 4 contained almost 30% of the data set, while flow 5 contained only about 4%. To complete the spatial anomaly detection step, the spatially anomalous trajectories detected by HDBSCAN were associated with the air traffic flow to which they were "closest" according to the weighted Euclidean distance from all identified air traffic flow centroids. This association led to the creation of the energy anomaly detection data sets.

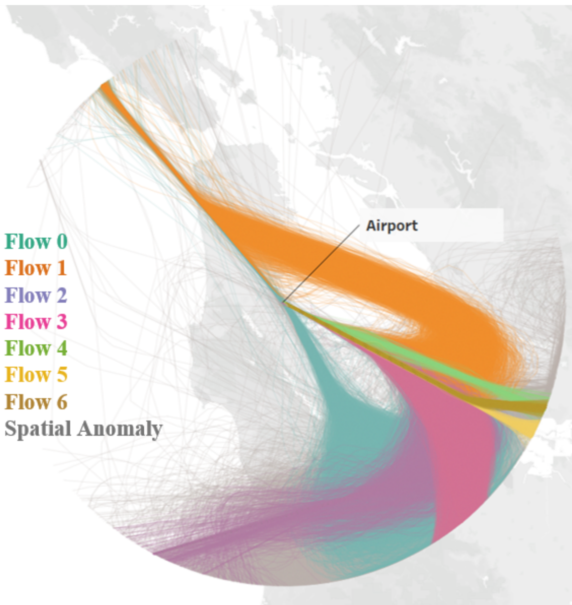

(a) Identified Air Traffic Flows

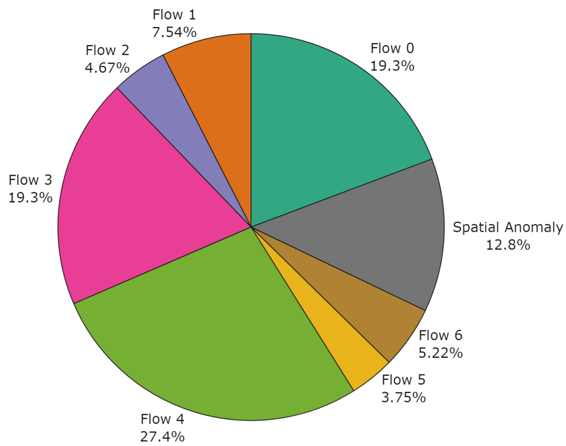

(b) Distribution of Trajectories

Figure 6: Spatial Anomaly Detection Results

5.4. *Energy Anomaly Detection*

The number of trajectories assigned to each energy anomaly detection data set as well as the portion of trajectories that had been detected as spatial anomalies associated with the data set are presented in Table 1. For each of the energy anomaly detection data sets, DBSCAN was applied following the same procedure.. The *sklearn* Python library's DBSCAN module [64] was leveraged to apply the DBSCAN algorithm. This library also provided the tools required for the z-normalization of the energy metrics, which was performed prior to applying DBSCAN to the energy anomaly detection data sets. As mentioned in Section 4, DBSCAN relies on two input parameters: $minSamples$ and ε, where, if $minSamples$ is fixed, ε may be varied to elicit a pre-determined fraction of outliers, or anomalies [44, 64]. The $minSamples$ parameter was set to be approximately 1% of the number of trajectories within each energy anomaly detection data set, which was consistent with the $minSamples$ setting in [46]. The exact value for each data set was dependent on the number trajectories associated with each identified air traffic flow in the spatial anomaly detection step. The ε parameter was varied until approximately 10% (between 9% and 11%) of the trajectories within the data set were detected as energy anomalies, which was consistent with the percentage of anomalies identified in other studies [16, 22].

Data Set	# Total Trajectories	% Spatial Anomalies
Flow 0	12,764	4.8
Flow 1	4,953	4.1
Flow 2	5,078	42.0
Flow 3	15,217	20.2
Flow 4	18,059	4.2
Flow 5	3,075	23.2
Flow 6	3,877	15.1

Table 1: Composition of Energy Anomaly Detection Data Sets

5.5. *Spatial and Energy Anomaly Relationship Exploration*

Of the most interest in exploring the relationship between spatial and energy anomalies is assessing the implication, if any, that detecting a trajectory as a spatial anomaly has on the subsequent energy anomaly detection. The utilization of DBSCAN to detect energy anomalies within each of the data sets led to each trajectory being assigned to one of the four categories introduced in Section 4: nominal (N), only spatial anomaly (S), only energy anomaly (E), both spatial and energy anomaly (B). The proportion of trajectories assigned to each category, for each energy anomaly detection data set, is displayed in Figure 7. It is evident these proportions vary between data sets; however, the sum of the proportions of the both energy and spatial (B) and only energy anomalies (E) are constant and equal to approximately 0.10.

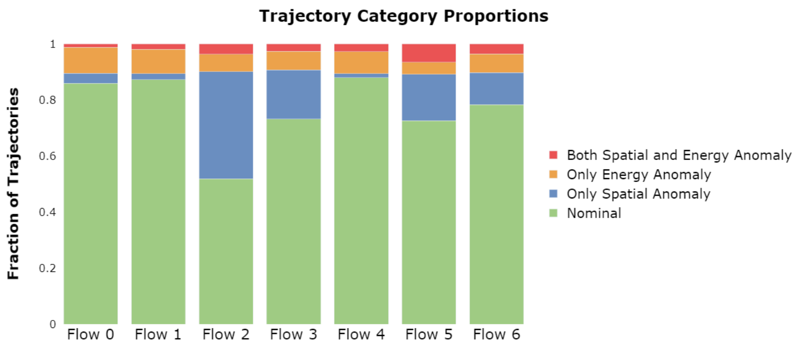

Figure 7: Trajectory Category Proportions within Each Energy Anomaly Detection Data Set

For each energy anomaly detection data set, the fraction of spatially nominal trajectories detected as energy anomalies and fraction of spatially anomalous trajectories detected as energy anomalies was computed, as illustrated in Figure 8. For example, for the flow 4 data set, only 8.2% of spatially nominal trajectories were detected as energy anomalies, whereas a much higher 65.6% of spatially anomalous trajectories were detected as energy anomalies. In

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 17 February 2021 doi:10.20944/preprints202102.0396.v1

general, for all flows with the exception of flow 2, the fraction of energy anomalies detected is significantly higher in the group of spatially anomalous trajectories than it is in the group of spatially nominal trajectories. The composition of the flow 2 data set was further investigated to understand the unique results.

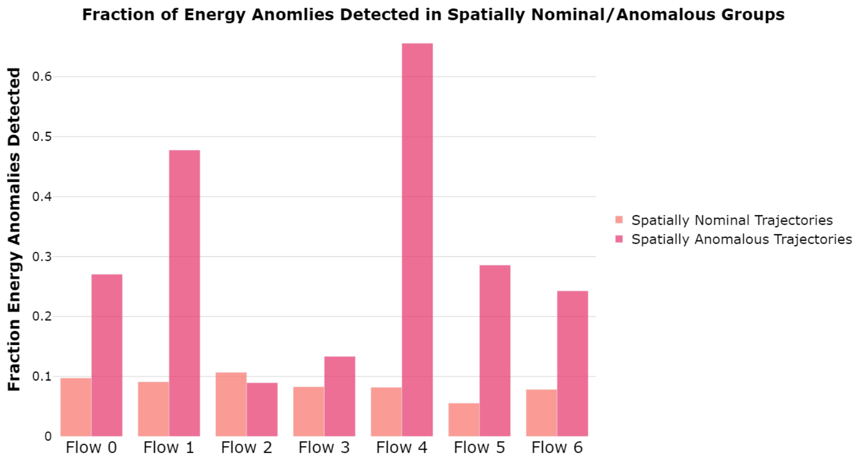

Figure 8: Implication of Spatial Anomalies

It is observed in both Table 1 and Figure 7 that the flow 2 data set consists of an unusually large portion of spatial anomalies (42.0%). It appeared that the majority of the spatial anomalies detected that are associated with flow 2 *should* have been assigned by the spatial anomaly detection algorithm to actually belong to flow 2 (rather than be anomalies). The data sets associated with the other identified air traffic flows were also examined and were not found to have similar issues. Due to a large majority of the *detected* spatial anomalies that do not appear to be *actual* spatial anomalies within the flow 2 data set, results were skewed and conclusions were difficult to draw. Because many of the *detected* spatial anomalies appear to have actually belonged in the spatially nominal trajectories group, this likely drove up the fraction of spatially nominal trajectories detected as energy anomalies and drove down the fraction of spatially anomalous trajectories detected as energy anomalies. Thus, this explained the uniqueness of the results for the flow 2 data set.

Figure 8 establishes that, in general, spatially anomalous trajectories are more likely to be detected as being energy anomalies than spatially nominal trajectories. While the identification of this relationship is important, it is equally important to note that not all spatial anomalies are subsequently detected as being energy anomalies. Table 2 displays the *likelihood ratio*, or ratio between the fraction of energy anomalies detected within the grouping of spatially anomalous trajectories and fraction of energy anomalies detected within the grouping of spatially nominal trajectories. This metric quantitatively indicates how much more or less likely it was for a trajectory to be detected as an energy anomaly given that the trajectory was detected as being spatially anomalous. For instance, for the previously mentioned flow 4 data set, the likelihood ratio was computed as 8.01, which indicates that spatially anomalous trajectories associated with flow 4 are about eight times more likely to be detected as energy anomalous than the spatially nominal trajectories. On average across all flows, excluding the flow 2 data set due to skewed results, spatially anomalous trajectories are between 1.6 to 8 times more likely to be identified as energy anomalies than spatially nominal trajectories. This result validates the first hypothesis presented in Section 4. Therefore, ATC and other operators should be aware of this difference and potential for off-nominal energy management to arise when an aircraft is following an off-nominal spatial path, whether this reason is due to ATC instructions or off-nominal pilot actions.

To continue with the analysis of the relationship and interdependencies between spatial and energy anomalies, the anomaly scores for each trajectory were computed, where the *sklearn* Python library's function to obtain pairwise distances [64] was utilized for this computation. The computation of the anomaly scores enabled the second hypothesis presented in Section 4 to be tested. The mean of the anomaly scores for each trajectory category was computed. The ratios of the mean anomaly score of each of the three anomalous categories of trajectories (S, E, and B) to the nominal category of trajectories were computed. These ratios are displayed in Figure 9. It is observed that, with the exception

Data Set	Likelihood Ratio	(S Mean)/(N Mean)	Welch's t-test p-value
Flow 0	2.77	1.02	0.000
Flow 1	5.24	1.02	0.024
Flow 2	0.84	0.99	1.000
Flow 3	1.61	1.02	0.000
Flow 4	8.01	1.03	0.001
Flow 5	5.15	1.02	0.000
Flow 6	3.10	1.04	0.000

Table 2: Spatial and Energy Anomaly Detection Relationship Exploration

of the flow 2 data set, the ratio of the mean anomaly score for only spatial anomalies (S) to the mean anomaly score for nominal trajectories (N) was greater than one. This indicates that, on average, the only spatial anomalies (S) had higher anomaly scores than the nominal trajectories (N), despite having been detected as being energy-nominal. Welch's one-sided t-test is performed to determine if this relationship is statistically significant. The *scipy* Python library [65] was leveraged to perform the Welch's one-sided t-test introduced in Section 4. The Welch's one-sided t-test performed tests for the null hypothesis that the two samples (S and N) have identical mean values [65]. The returned p-values for this statistical test are displayed in Table 2. It is noted that, again, with the exception of the flow 2 data set, the p-values returned were less than the commonly utilized 0.05 statistical significance threshold, such that the null hypothesis (the only spatial (S) anomaly score mean is greater than the nominal (N) score mean) is rejected. Therefore, the second hypothesis presented in Section 4 is validated.

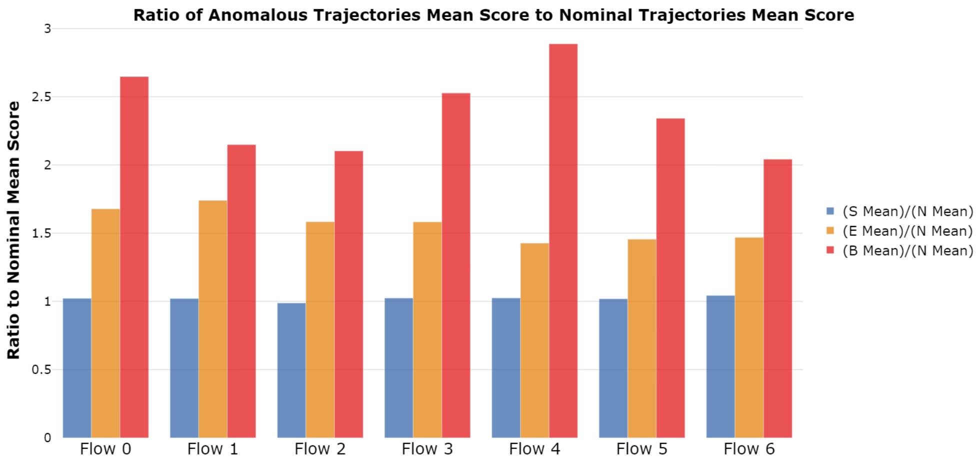

Figure 9: Energy Anomaly Scores Ratios

Figure 9 also indicates that both spatial and energy anomalies appeared to be more "severe" in that their anomaly scores appear to be greater than the only energy anomaly (E) anomaly scores. Figure 10 displays the distributions of anomaly scores, grouped by trajectory category, for all trajectories within all flows. Through visual inspection, it is evident the anomaly scores of the only energy anomalies and both spatial and energy anomalies are different. The distribution of the anomaly scores for the both spatial and energy anomalies is skewed toward much higher anomaly scores. Breaking the distributions down by individual flows produced similar results. Therefore, the final hypothesis presented in Section 4 is validated. This indicates that the underlying mechanisms may be dissimilar for only energy anomalies (E) and both spatial and energy anomalies (B). Further, it appeared that more "severe" energy anomalies occurred if an aircraft was observed to experience off-nominal spatial states, which is of interest to aviation operators to consider to facilitate safe and efficient ATM.

Considering all metrics discussed thus far, the data set associated with flow 5 produced metrics that are generally

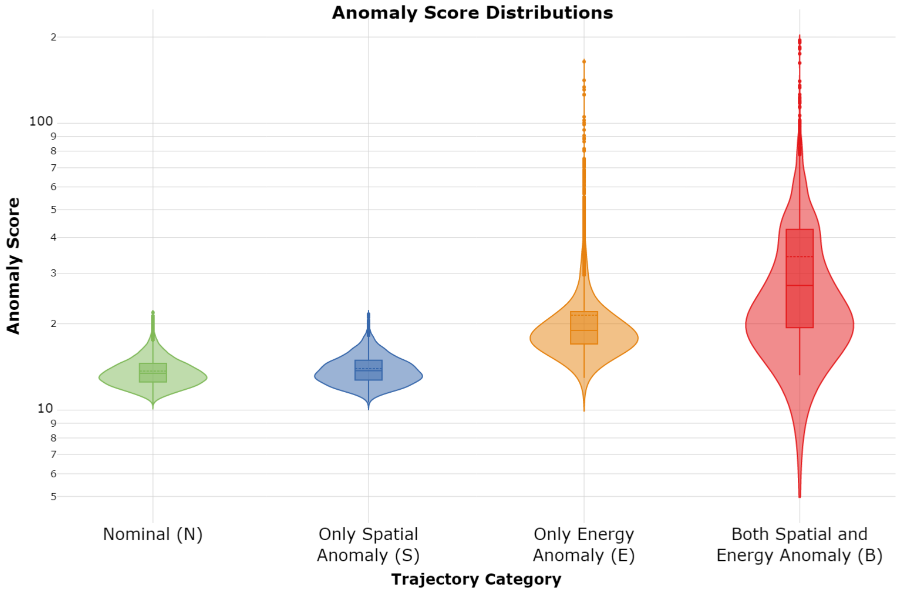

Figure 10: Trajectory Category Energy Anomaly Score Distribution

in the “middle” of the metric ranges. As such, this data set is selected to present a more in-depth analysis. To further aid in the investigation of the relationship between only energy anomalies (E) and both spatial and energy anomalies (B), the t-distributed stochastic neighbor embedding (t-SNE) [66] dimensionality reduction technique was leveraged for visualization of the high-dimensional feature vectors in two dimensions. Figure 11 displays the two components resulting from applying t-SNE, colored by whether the trajectory was detected as energy nominal (nominal and only spatial anomaly categories), an only energy anomaly, or both spatial and energy anomaly. The reduced-dimension representation implies a separation between only energy anomalies and both spatial and energy anomalies, where the two categories are generally clustered at two “ends” of the resulting blob of the trajectory representations. ATC and other operators should be aware of the underlying differences in energy metrics sequences between trajectories experiencing poor energy management while remaining on a nominal spatial path versus veering off a nominal spatial path to ensure safe and efficient operations within the terminal airspace.

Individual trajectories assigned to the three anomaly categories within the flow 5 data set are presented. The spatial profiles of these trajectories, overlaid on the profiles of all spatially nominal trajectories, are displayed in Figure 12. Similarly, the energy profiles of these three trajectories in the three energy metrics of interest (SPE, SKE, STER), overlaid on the bounds of energy-nominal trajectories, are displayed in Figure 13. It is evident that the selected only spatial anomaly’s trajectory (in blue) deviates significantly from the spatially nominal trajectories (in green). However, despite the significant spatial deviation, the energy profile remains nominal throughout the entire approach, which supports the observation that a spatial anomaly is not always also an energy anomaly. The complement is true when assessing the selected only energy anomaly’s trajectory. This trajectory (in orange) is in near exact alignment with other spatially nominal trajectories, though still experiences large deviations in its energy metrics. Finally, it is evident that the selected both spatial and energy anomaly’s spatial and energy profiles deviate significantly, and that, in this case, the deviations seem more severe than the selected only spatial and only energy anomalies, in the respective anomalous dimension.

6. Conclusions

This paper introduced the concepts of spatial and energy anomalies and presented a novel, unified framework for the detection of spatial and energy anomalies in ADS-B trajectory data (and associated metrics). Such framework, which currently does not exist in the literature, consists of three main parts - a data processing procedure, a spatial

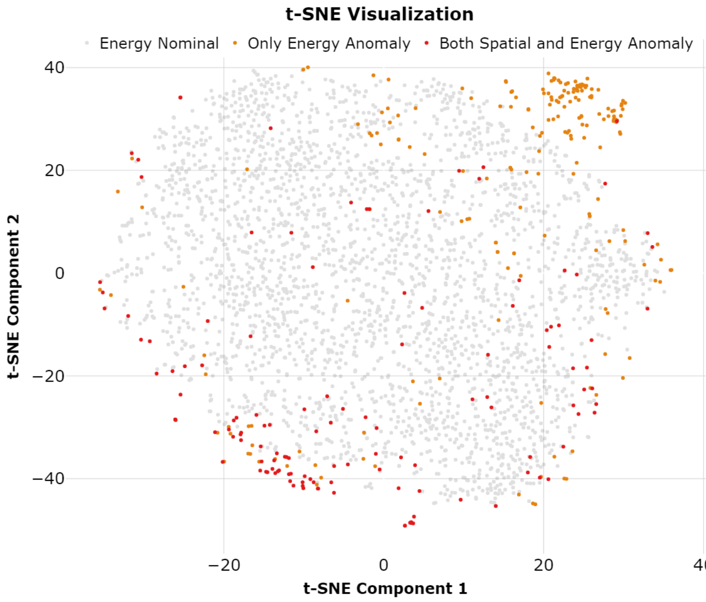

Figure 11: Comparison of the Appearance of Only Energy and Both Spatial and Energy Categories in Two Dimensions

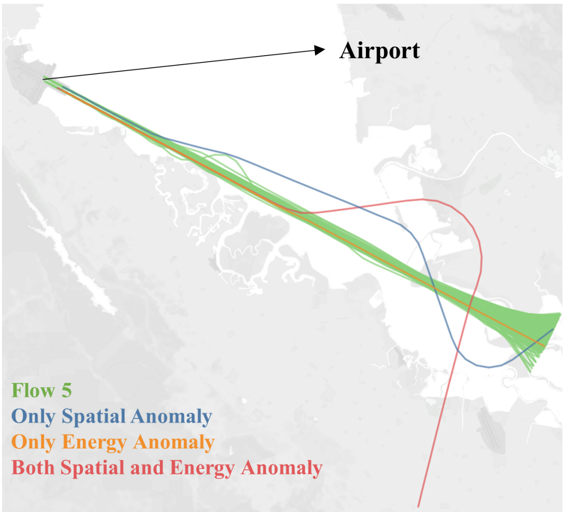

Figure 12: Flow 5, Spatial Profiles and Selected Anomalies

anomaly detection method, and an energy anomaly detection method. Once developed and implemented, this framework was used to explore the relationship and interdependencies between spatial and energy anomalies, which may be of great interest to aviation operators and ATM decision-makers.

For the purpose of this paper, trajectories available from the OpenSky network's historical database of flights arriving at San Francisco International Airport between June 1^{st}, 2019 and September 30^{th}, 2019 were considered. Both the spatial and energy anomaly detection methods used to demonstrate this framework are clustering-based methods, where HDBSCAN with a weighted Euclidean distance is used for spatial anomaly detection and DBSCAN is used for energy anomaly detection. While these two methods were chosen, there is no restriction on the type of anomaly detection method that can be used for either spatial or energy anomaly detection when implementing this framework.

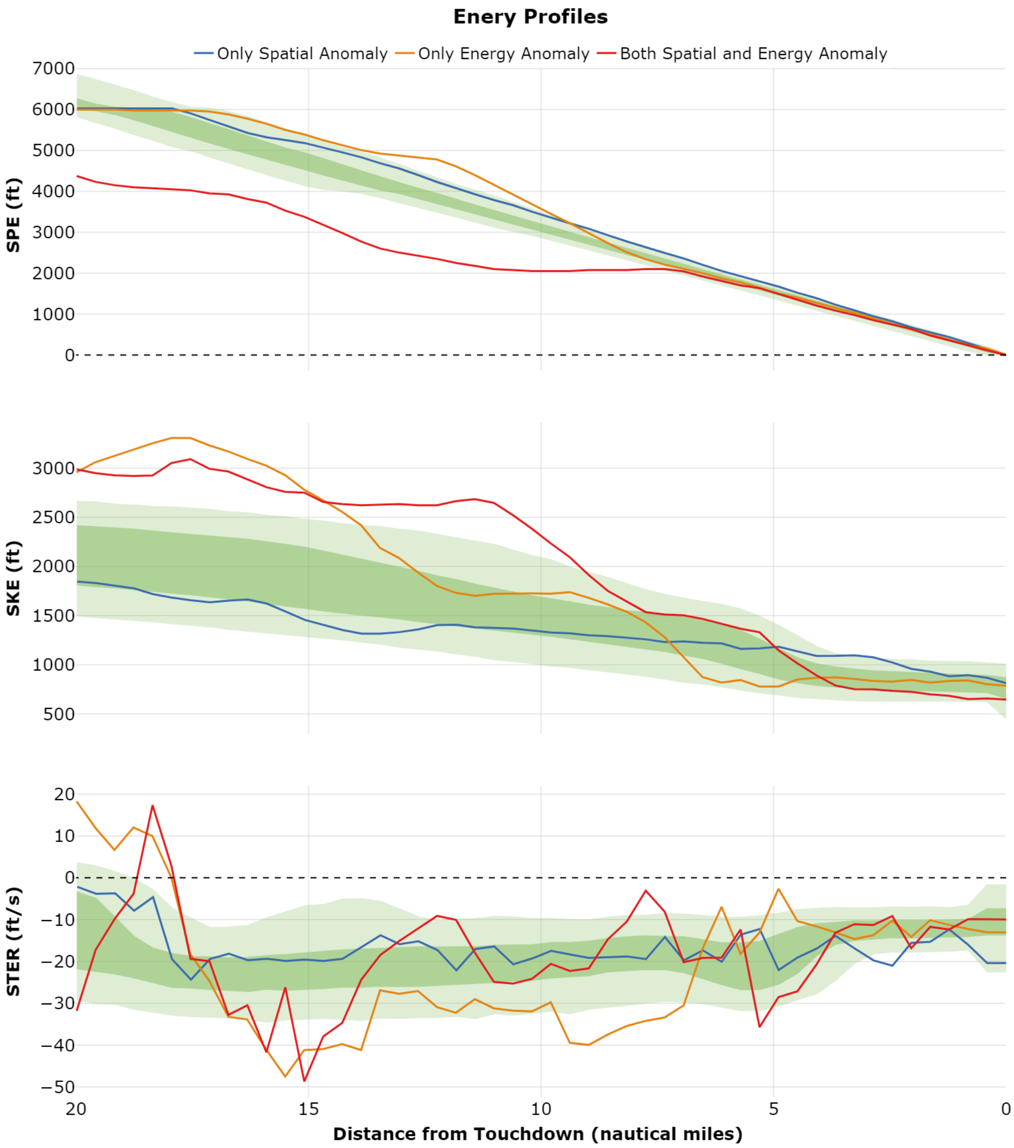

Figure 13: Flow 5, Energy Profiles and Selected Anomalies

Exploration of the relationship between spatial and energy anomalies revealed that, in general, if a trajectory is detected as being a spatial anomaly, it is more likely to be detected as also being an energy anomaly. Whether a flight's trajectory is not conforming to an identified set of air traffic flows representing standard spatial operations due to pilot actions or instructions from ATC, it is appropriate to make pilots aware of increased emphasis on the importance of managing the aircraft's energy state such that potentially risky or unsafe situations do not arise. Further, the average anomaly score for trajectories detected as only spatial anomalies is generally higher than the average anomaly score for nominal trajectories. Aviation operators, such as air traffic controllers, may benefit from this observation, as it may factor into decision-making in instances where there is the potential to instruct an aircraft to spatially deviate from standard operations. Additionally, this research revealed that it appears to be underlying differences in energy metrics sequences of only energy anomalies and both spatial and energy anomalies, which should be considered when making decisions on actions that air traffic controllers or pilots should take.

The primary benefit of this framework is the ability to analyze aircraft trajectories for both spatial and energy metrics conformance and investigate the relationship between the spatial and energy anomalies. In addition, this framework provides aviation operators with the ability to down-select specific categories of anomalies for further analysis by subject matter experts. In future work, the application of this framework will be extended to include flights departing within the terminal airspace. Due to differences between arrival and departure procedures, the relationship between spatial and energy anomalies cannot be extrapolated to departures.

References

[1] Federal Aviation Administration, FAA Aerospace Forecast Fiscal Years 2020-2040, https://www.faa.gov/data_research/aviation/aerospace_forecasts/media/FY2020-40_FAA_Aerospace_Forecast.pdf, 2020.
[2] Federal Aviation Administration, Modernization of the U.S. Airspace, https://www.faa.gov/nextgen/, 2020.
[3] Single European Sky (SES), Single European Sky, https://www.sesarju.ed/background-ses, 2020.
[4] Puranik, T.G., A Methodology for Quantitative Data-driven Safety Assessment for General Aviation, Ph.D. thesis, Georgia Institute of Technology, 2018.
[5] National Aeronautics and Space Administration, Strategic Implementation Plan 2019 Update, 2019.
[6] National Academies of Sciences, Engineering, and Medicine, In-Time Aviation Safety Management: Challenges and Research for an Evolving Aviation System (2018), The National Academies Press, Washington, DC, United States, 2018.
[7] D. Rinehart, H. Jimenez, M. Blake, Development of real-time system-wide safety assurance definitions and concept fundamentals, in: 16th AIAA Aviation Technology, Integration, and Operations Conference, AIAA, Washington, DC, United States, pp. 1–12.
[8] E. Mangortey, J. Gilleron, G. Dard, O. J. Pinon, D. N. Mavris, Development of a data fusion framework to support the analysis of aviation big data, in: AIAA SciTech 2019 SciTech Forum, AIAA, pp. 1–17.
[9] X. Olive, J. Morio, Trajectory clustering of air traffic flows around airports, Aerospace Science and Technology 84 (2019) 776–781.
[10] M. C. R. Murca, R. J. Hansman, L. Li, P. Ren, Flight trajectory data analytics for characterization f air traffic flows: A comparative analysis of terminal area operations between new york, hong kong and sao paulo, Transportation Research Part C 97 (2018) 324–347.
[11] R. Deshmukh, I. Hwang, Anomaly detection using temporal logic based learning for terminal airspace operations, in: AIAA SciTech 2019 SciTech Forum, AIAA, pp. 1–11.
[12] S. Pozzi, C. Valbonesi, V. Beato, R. Volpini, F. M. Giustizieri, F. Lieutaud, A. Lieu, Safety monitoring in the age of big data, in: Ninth USA/Europe Air Traffic Management Research and Development Seminar (ATM2011), pp. 1–9.
[13] A. Gavrilovski, H. Jimenez, D. N. Mavris, A. H. Rao, S. Shin, I. Hwang, K. Marais, Challenges and opportunities in flight data mining: A review of the state of the art, in: AIAA Infotech@Aerospace, San Diego, California, United States, pp. 1–18.
[14] B. Matthews, S. Das, K. Bhaduri, K. Das, R. Martin, N. Oza, Discovering anomalous aviation safety events using scalable data mining algorithms, Journal of Aerospace Information Systems 10 (2013).
[15] M. Sharma, K. Das, M. Bilgic, B. Matthews, D. Nielsen, N. Oza, Active learning with rationales for identifying operationally significant anomalies in aviation, in: Joint European Conference on Machine Learning and Knowledge Discovery in Databases, pp. 209–225.
[16] S. Das, B. L. Matthews, A. N. Srivastava, N. C. Oza, Multiple kernel learning for heterogeneous anomaly detection: Algorithm and aviation safety case study, in: Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '10, Association for Computing Machinery, New York, NY, USA, 2010, p. 47–56.
[17] L. Li, S. Das, J. R. Hansman, R. Palacios, A. N. Srivastava, Analysis of flight data using clustering techniques for detecting abnormal operations, Journal of Aerospace Information Systems 12 (2015) 587–598.
[18] L. Basora, T. D. Xavier Olive, Recent advances in anomaly detection methods applied to aviation, Aerospace 6 (2019).
[19] T. G. Puranik, D. N. Mavris, Anomaly detection in general aviation operations using energy metrics and flight-data records, Journal of Aerospace Information Systems 15 (2018) 22–35.
[20] T. Puranik, D. N. Mavris, Identification of instantaneous anomalies in general aviation operations using energy metrics, Journal of Aerospace Information Systems 17 (2020) 51–65.
[21] L. Li, J. Hansman, R. Palacios, R. Welsch, Anomaly detection via a gaussian mixture model for flight operation and safety monitoring, Reliability Engineering & System Safety 193 (2020).
[22] K. Sheridan, T. Puranik, E. Mangortey, O. J. Pinon, M. Kirby, D. N. Mavris, An application of dbscan clustering for flight anomaly detection during the approach phase, in: 2020 AIAA SciTech Forum, AIAA, pp. 1–20.
[23] Mayara Conde Rocha Murca, Data-Driven Modeling of Air Traffic Flows for Advanced Air Traffic Management, Ph.D. thesis, Massachusetts Institute of Technology, 2018.
[24] Airbus, Accidents by Flight Phase, https://accidentstats.airbus.com/statistics/accident-by-flight-phase, 2020.
[25] F. Jackman, Nearly half of commerical jet accidents occur during final approach, landing, 2014.
[26] B. Sridhar, Applications of machine learning techniques to aviation operations: Promises and challenges, in: 2020 International Conference on Artificial Intelligence and Data Analytics for Air Transportation (AIDA-AT), pp. 1–12.
[27] G. Jarry, D. Delahaye, F. Nicol, E. Feron, Aircraft atypical approach detection using functional principal component analysis, Journal of Air Transport Management 84 (2020).
[28] T. Puranik, H. Jimenez, D. Mavris, Energy-based metrics for safety analysis of general aviation operations, Journal of Aircraft 54 (2017) 2285–2297.
[29] M. Gariel, A. N. Srivastava, E. Feron, Trajectory clustering and an application to airspace monitoring, IEEE Transactions on Intelligent Transportation Systems 12 (2011) 1511–1524.
[30] R. Deshmukh, I. Hwang, Incremental-learning-based unsupervsied anomaly detection algorithm for terminal airspace operations, Journal of Aerospace Information Systems 19 (2019) 362–384.

[31] R. Deshmukh, D. Sun, I. Hwang, Data-driven precursor detection algorithm for terminal airspace operations, in: Thirteenth USA/Europe Air Traffic Management Research and Development Seminar, pp. 1–7.
[32] Raj Deshmukh, Data-Driven Anomaly and Precursor Detection in Metroplex Airspace Operations, Ph.D. thesis, Purdue University, 2020.
[33] K. Kim, I. Hwang, Terminal airspace anomaly detection using temporal logic learning, in: International Conference on Research in Air Traffic 2018, ICRAT, pp. 1–4.
[34] X. Olive, J. Grignard, T. Dubot, J. Saint-Lot, Detecting controllers' actions in past mode s data by autoencoder-based anomaly detection, in: SESAR Innovation Days, pp. 1–8.
[35] X. Olive, L. Basora, Identifying anomalies in past en-route trajectories with clustering and anomaly detection methods, in: Thirteenth USA/Europe Air Traffic Management Research and Development Seminar, pp. 1–10.
[36] X. Olive, L. Basora, B. Viry, R. Alliger, Deep trajectory clustering with autoencoders, in: International Conference on Research in Air Transportation 2020, ICRAT, pp. 1–8.
[37] X. Olive, L. Basora, Detection and identification of significant events in historical aircraft trajectory data, Transportation Research Part C: Emerging Technologies 119 (2020).
[38] A. Tanner, M. Strohmeier, Anomalies in the sky: Experiments with traffic densities and airport runway use, in: 7th OpenSky Workshop 2019, pp. 1–12.
[39] S. Kaushik, An introduction to clustering and different methods of clustering, https://www.analyticsvidhya.com/blog/2016/11/an-introduction-to-clustering-and-different-methods-of-clustering/, 2016.
[40] M. Enriquez, Identifying temporally persistent flows in the terminal airspace via spectral clustering, in: Air Traffic Management R&D Seminar, pp. 1–8.
[41] A. Eckstein, Automated flight track taxonomy for measuring benefits from performance based navigation, in: 2009 Integrated Communications, Navigation and Surveillance Conference, pp. 1–12.
[42] F. Rehm, Clustering of flight tracks, in: AIAA Infotech@Aerospace, pp. 1–9.
[43] M. Enriquez, C. Kurcz, A simple and robust flow detection algorithm based on spectral clustering, in: International Conference on Research in Air Transportation, Berkley, California, United States, pp. 1–6.
[44] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, A density-based algorithm for discovering clusters in large spatial databases with noise, in: Second International Conference on Knowledge Discovery and Data Mining, AAAI Press, 1996, p. 226–231.
[45] X. Olive, L. Basora, A python toolbox for processing air traffic data: A use case with trajectory clustering, in: 7th OpenSky Workshop 2019, pp. 1–12.
[46] L. Basora, J. Morio, C. Mailhot, A trajectory clustering framework to analyse air traffic flows, in: 7th SESAR Innovation Days, Belgrade, Serbia, pp. 1–8.
[47] P. Besse, B. Guillouet, J.-M. Loubes, F. Royer, Review & perspective for distance based clustering of vehicle trajectories, IEEE Transactions on Intelligent Transportation Systems 17 (2016) 3306–3317.
[48] S. Corrado, T. Puranik, O. Pinon, D. Mavris, Trajectory clustering within the terminal airspace utilizing a weighted distance function, in: Proceedings of the 8th OpenSky Symposium 2020, pp. 1–10.
[49] S. R. Proud, Go-around detection using crowd-sourced ads-b position data, Aerospace 7 (2019).
[50] B. Figuet, R. Monstein, M. Waltert, S. Barry, Predicting airplane go-arounds using machine learning and open-source data, in: Proceedings of the 8th OpenSky Symposium 2020, pp. 1–12.
[51] J. Bro, Fdm machine learning: An investigation into the utility of neural networks as a predictive analytic tool for go around decision making, Journal of Applied Sciences and Arts 1 (2017).
[52] L. Li, M. Gariel, R. J. Hansman, R. Palacios, Anomaly detection in onboard-recorded flight data using cluster analysis, in: 2011 IEEE/AIAA 30th Digital Avionics Systems Conference, pp. 4A4–1–4A4–11.
[53] T. G. Puranik, H. Jimenez, D. N. Mavris, Utilizing energy metrics and clustering techniques to identify anomalous general aviation operations, in: AIAA SciTech Forum, AIAA, pp. 1–19.
[54] T. G. Puranik, D. N. Mavris, Identifying instantaneous anomalies in general aviation operations, in: AIAA Aviation Forum, AIAA, pp. 1–15.
[55] G. Jarry, D. Delahaye, E. Feron, Trajectory approach analysis: A post-operational aircraft approach analysis tool, in: 9th SESAR Innovation Days, Athens, Greexe, pp. 1–8.
[56] R. J. G. B. Campello, D. Moulavi, J. Sander, Density-based clustering based on hierarchical estimates, in: Advances in Knowledge Discovery and Data Mining, AAAI Press, 2013, pp. 160–172.
[57] A. Fernandez, D. Martinez, P. Hernandez, S. Cristobal, F. Schwaiger, J. M. Nunez, Flight data monitoring (fdm) unknown hazards detection during approach phase using clustering techniques and autoencoders, in: SESAR Innovation Days 2019, pp. 1–8.
[58] L. McInnes, J. Healy, S. Astels, hdbscan: Hierarchical density based clustering, The Journal of Open Source Software 2 (2017).
[59] B. L. Welch, The generalization of 'student's' problem when several different population variances are involved, Biometrika 34 (1947) 28–35.
[60] M. Schäfer, M. Strohmeier, V. Lenders, I. Martinovic, M. Wilhelm, Bringing up opensky: A large-scale ads-b sensor network for research, in: 13th IEEE/ACM International Symposium on Information Processing in Sensor Networks (IPNS), pp. 38–94.
[61] X. Olive, Traffic, a toolbox for processing and analysing air traffic data, Journal of Open Source Software (2019) 1518.
[62] Iowa State University, Iowa Environmental Mesonet, https://mesonet.agron.iastate.edu/request/download.phtml, 2021.
[63] T. Blajev, W. Curtis, Go-Around Decision-Making and Execution Project, 2017.
[64] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine learning in Python, Journal of Machine Learning Research 12 (2011) 2825–2830.
[65] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, SciPy 1.0 Contributors, SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, Nature

Methods 17 (2020) 261–272.
[66] L. van der Maaten, G. Hinton, Visualizing data using t-sne, Journal of Machine Learning Research 9 (2008) 2579–2605.