Identifying Anomalies in past en-route Trajectories with Clustering and Anomaly Detection Methods

Xavier Olive, Luis Basora
ONERA DTIS, Université de Toulouse
Toulouse, France

Abstract—This paper presents a framework to identify and characterise anomalies in past en-route Mode S trajectories. The technique builds upon two previous contributions introduced in 2018: it combines a trajectory-clustering method to obtain the main flows in an airspace with autoencoding artificial neural networks to perform anomaly detection in flown trajectories. The combination of these two well-known Machine Learning techniques (ML) provides a useful reading grid associating cluster analysis with quantified level of abnormality.

The methodology is applied to a sector of the French Bordeaux Area Control Center (ACC) during its 385 hours of operation over seven months of ADS-B traffic. The results provide a good taxonomy of deconfliction measures and weather-related ATC actions. The application of this work is manyfold, ranging from safety studies estimating risks of midair collision, to complexity and workload assessments of traffic when a sector is operated, or to the constitution of a database of ATC actions ensuring aircraft separation. This database could be used to train further ML techniques aimed at improving the state of the art of deconfliction algorithms.

Keywords—trajectory analysis; trajectory clustering; anomaly detection; autoencoders

I. INTRODUCTION

The forecast growth of air traffic will lead to an increase in safety issues especially in the already congested high-density airspaces of the US and Europe. The respective air transport authorities are both addressing the problem in part by putting into place initiatives to better collect and analyze recorded flight information such as the Flight Operational Quality Assurance (FOQA) data programs. On the other hand, the availability of vast amounts of data via open or public sources like ADS-B along with the advances in the Machine Learning (ML) field smooth the way towards an automatic discovery of unsafe flight patterns and other operational anomalies.

This paper presents a methodology to identify operationally significant events in recorded ADS-B flight data which can be associated with Air Traffic Control (ATC) actions in en-route traffic. Controllers work to ensure safety and optimize traffic flows; they provide various clearances and take conflict resolution measures. The identification of traffic situations and associated ATC actions could address many applications, including safety analyses or preparations of ATC training simulations. Identifying situations which could have led to a separation loss, analyzing context and proximate events are valuable inputs for Collision Risk Models (CRM) [3], resulting in estimations of mid-air collision risks. Locating hot spots based on identified situations could also help assess the complexity of traffic inside a given sector when it is in operation.

The detection of significant events can be achieved by automatically identifying flights presenting a certain degree of anomaly. The goal of our research is to design a methodology for finding anomalies in the en-route phase of the flights in a given airspace or sector. Our approach combines a trajectory clustering method [1] to identify air traffic flows within the considered airspace with a technique to detect anomalies [2] in each flow. The proposed method highlights trajectories, which in spite of not being considered directly as outliers by the clustering algorithm, still present a degree of dissimilarity significant enough to justify further analysis. The final goal would be to understand the causes and implications of what could be labelled as potential ATC actions.

The evaluation of our methodology is performed on a seven months' traffic data from a sector in the Bordeaux Area Control Center (ACC) for periods when the sector was operationally deployed according to the Sector Configuration Plans (SCP), also known as *opening schemes*. Our method has been able to identify many ATC situations of various types, to be related to the structure of the traffic in the sector.

The rest of the paper is organized as follows. Section II reviews the state of the art with respect to clustering and anomaly detection techniques. Section III presents the considered use case over LFBBPT sector controlled by Bordeaux ACC. Section IV recalls the various steps of clustering and anomaly detection that are applied to the considered dataset. Then, Section V presents the results and Section VI analyses specific situations found by our toolchain. In conclusion, we recall the most significant benefits of our approach, main results and the potential impact of our contribution.

II. LITERATURE REVIEW AND RELATED WORK

Recently, Machine Learning (ML) techniques have brought a new perspective to ATM related problems: unsupervised machine learning mainly addresses the description of data and of its underlying structure: clustering [1], [4], [5], [6], and anomaly detection [2], [7], [8] fall into this category; supervised machine learning applies to labeled datasets and focuses on predicting aircraft trajectories [9], [10], [11]. Reinforcement learning addresses more traditional optimisation problems [12] where the search space may be explored through interactions with a simulator or based on a history of situations.

Bundling [13], [14] is a reliable technique related to flow identification. It is a powerful visualisation tool which distorts and groups pieces of trajectories; this method is well suited to exploration analytics but fails to provide a label to each trajectory, essential for grouping samples into categories to be used for further studies.

A natural way to identify trajectory flows is by using clustering techniques. Such algorithms group similar trajectories together in clusters, where each cluster is associated with a traffic flow represented by a *centroid*, i.e. the mean of the trajectories in the cluster. A number of methods in the literature [5], [6], [15], [16], [17], [18] exist to cluster flight trajectories into flows, most of them based on density-based clustering algorithms such as DBSCAN [19].

Basora et al. [1] describe a clustering method specifically designed to identify the traffic flows in a sector, based on a progressive clustering technique originally developed by Andrienko et al. [16], [17]. This method was successfully applied to identify the flows for all the sectors in the Bordeaux Area Control Center (ACC) by selecting a set of six clustering variables.

Analysing outliers resulting from a clustering algorithm can be a relevant method to find anomalies in a dataset. Olive [6] proposed a different technique for identifying converging flows in the terminal area of Toulouse which helps understand how approaches are managed. The analysis of outliers provides elements to understand and assess specific situations calling for more in-depth safety analyses.

Other techniques for anomaly detection have been proposed in recent years: Das et al. [7] introduced a technique called Multi-Kernel Anomaly Detection (MKAD) based on a One-class Support Vector Machine (SVM), which has been applied to FOQA data to detect significant events in the approach phase [20], [21], [22]. The main advantage of this method is the ability to consider continuous and discrete variables through a linear combination of many kernel functions for each variable type. On the other hand, clustering based anomaly detection (ClusterAD) [23], [24] associates nominal flights with clusters and considers the remaining flights not belonging to a specific cluster (outliers) as anomalous flights.

In spite of the good quality results, limitations in both ClusterAD and MKAD have been pointed out in studies [25], [26]: the authors pointed out the need for dimensionality reduction, the poor sensitivity to short duration anomalies, and the inability to detect anomalies in latent features. They proposed different approaches [26], [27] based on a Vector Auto-Regressive (VAR) technique to model the FOQA data and identify anomalies in the flights.

Alternatively, Nanduri et al. [25] describe the application of Recurrent Neural Networks (RNNs) to overcome the limitations enumerated above. In particular, they showed that a Long-Short Term Memory (LSTM) architecture can perform better than MKAD in detecting flight anomalies. The authors trained several RNN architectures to detect 11 canonical anomalies found in the literature [7], [23], [28] by generating data in the format of FOQA from an X-Plane simulation. As

a result, RNNs were able to identify 8 out of 11 anomalies, whereas MKAD was only able to identify 6 of them.

In this paper, we focus on autoencoders, a particular kind of neural network which have recently proved successful at anomaly detection. They have already been used to find breakpoints in time series [29], to predict realistic transitions in sector configurations [30] and to detect and distinguish atypical situations (mostly weather related) and controllers' actions in Mode S data [2]. Autoencoders are comparable to compression methods: they are trained to reconstruct, i.e. compress then decompress, data (trajectories) passed in input. In practice, they learn to reconstruct most samples in a training dataset and fail to reconstruct the more atypical ones. Anomaly detection is based on the distance, the *reconstruction error*, between input trajectories and their reconstructed copies. In particular, the analysis of distributions of reconstruction errors has been a powerful asset in the context of ATM data.

III. SCENARIO AND DATASETS

Mode S has become one of the most important technologies in air traffic management as it supports the operation of secondary surveillance radar (SSR), traffic alert and collision avoidance systems (TCAS), and Automatic Dependent Surveillance–Broadcast (ADS-B). In practice, transponders in aircraft are selectively interrogated by sensors (radars) to provide situational awareness through the exchange of binary encoded information.

The OpenSky Network [31] is a crowd-sourced sensor network collecting such air traffic data. The collected data used for this study contains only ADS-B data of aircraft flying a specific sector in the French Bordeaux Area Control Center (ACC) between January 1st and August 6th 2017.

Each ACC, in charge of providing air traffic control services to controlled flights within its airspace, is subdivided into elementary sectors that are used or combined to build control sectors operated by a pair of air traffic controllers. Airspace sectorisation consists in partitioning the overall ACC airspace into a given number of these control sectors. In most centres, the set of control sectors deployed, i.e. the *sector configuration*, varies throughout the day. Basically, sectors are split when controllers' workload increases, and merged when it decreases.

We call Sector Configuration Plan (SCP) the sequence of sector configurations to be deployed within an ACC throughout a day of operations. This plan, also known as *opening scheme*, is established by the Flow Management Position (FMP) on a daily basis and contains the information regarding the opening time slots for each sector configuration. The SCP dataset for this study was provided by the operational team in Bordeaux ACC in order to enable the selection of trajectories operated in each sector.

We have chosen for our study a specific sector, LFBBPT, and selected in the SCP dataset the time intervals when the sector was operationally deployed, meaning that a pair of controllers were in charge of the sector. We have considered all the sector opening slots of longer than 30 minutes and selected

all the associated trajectories (only for aircraft equipped with ADS-B capable transponders) crossing the sectors during these intervals. The goal of using the SCP is for the traffic under analysis to be representative of operational situations with a level of workload deemed acceptable by the controllers.


Fig. 1. Scenario with en-route traffic in Bordeaux ACC within LFBBPT sector. A traffic sample of 5000 trajectories is plotted for reference.

The scenario and the geographic footprint of the sector (available on the eAIP France, ENR 3.8[1]) are illustrated in Fig. 1 with a subsample of 5000 trajectories crossing the sector while it is in operation. Trajectories have been preprocessed to filter the noise out and extract the trajectory parts within the area of interest (sector LFBBPT, above 19,500 ft). Also, we have removed flights spending less than five minutes within the sector.

In order to meet the prerequisites of the clustering and anomaly detection techniques, the resulting trajectories have been resampled so as to have an equal number of points which we have fixed at 50 points per trajectory. During the 385 hours of operation of LFBBPT, a dataset of 14,461 trajectories has been used to evaluate our approach.

IV. METHODOLOGY

In this section, we describe the methodology for anomaly detection in the en-route traffic of an airspace. First, we identified the air traffic flows crossing the sector when operationally deployed. The flow identification is performed by the trajectory clustering algorithm developed in [1]. The output of the algorithm is the set of trajectory clusters and the set of outliers, i.e. the trajectories which could not be allocated to a cluster and which therefore should already present some degree of abnormality. The main goal in this phase is not so much to find anomalies in the traffic as to properly identify the air traffic flows. Then, we apply an autoencoder-based anomaly detection method [8] to detect events in the identified flows.

A. Flow Identification

In the clustering algorithm described in [1] a trajectory within a sector is represented with only its first and last points corresponding to the entry and exit points of the flight in the

---

[1]https://www.sia.aviation-civile.gouv.fr

sector. Therefore, for the clustering algorithm, the notion of similarity between trajectories is limited in this case to the way the flights enter and leave the sector. The authors found that this level of trajectory representation was sufficient to properly identify the flows as seen by the operational teams in the Bordeaux ACC sectors and so should be applicable to our scenario case.

The steps of the clustering method are illustrated in Fig. 2. First, a matrix is built from the trajectory dataset generated in the preprocessing phase, where each line of the matrix is a vector of six dimensions representing the trajectory entry/exit points to/from the sector, i.e. the two geographic coordinates and altitudes of each entry/exit point. A logarithmic function is applied to the altitude data, such that two stable (levelled off) cruise trajectories at high levels (e.g. FL330 and FL370) will be considered closer than two potentially evolving (stable, climbing or descending) trajectories interacting at lower levels (e.g. FL290 and FL330). This technique has been designed to help the algorithm separate the two types of trajectories. Also, all matrix values have been standardized to have zero mean and unit variance. The resulting matrix is the input to the DBSCAN algorithm which identifies the initial set of clusters (flows) and the outliers.
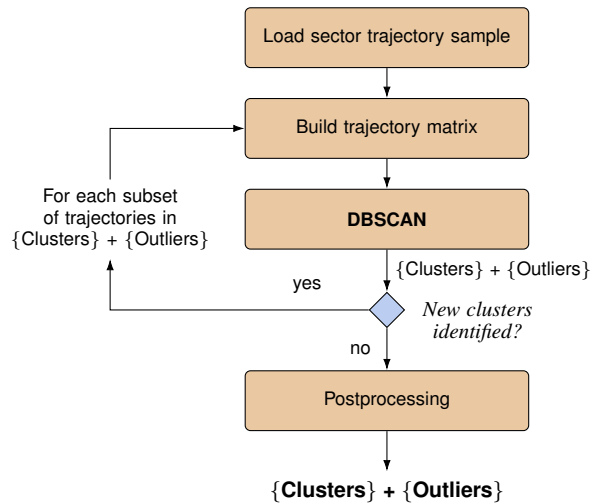

Fig. 2. Flow identification method.

As some of the flows were not always well separated after the first application of DBSCAN, we refined the clusters obtained by applying DBSCAN again separately to each of them, using a technique called *progressive clustering approach* [16], [17]. This refinement can be achieved thanks to the effect of data standardization which transforms data differently at the cluster and sub-cluster level. Moreover, we perform at every level of the recursion a fine-tuning of DBSCAN parameters, i.e. the density parameter ε (the maximum distance between two samples for them to be considered as in the same neighborhood) and the minimum sample size n, aimed at identifying additional flows within the clusters and outliers.

The iterative flow identification process stops when no new clusters are identified. Eventually, a post-processing phase is

initiated after the cluster centroids are computed. A fusion step was also added by the authors to compensate for the potential effect of over-clustering that can be caused by the iterative application of DBSCAN.

B. Autoencoders for Anomaly Detection

We implemented the method presented in [2] that had already been successfully applied to detect controllers' actions before aircraft enter the Terminal Maneuvering Areas of major airports. In this paper, we applied this technique en-route, on flows identified by trajectory clustering rather than on flows identified by city pairs. The technique uses autoencoders and work on the distribution of reconstruction errors.

1) Autoencoders: Autoencoders are artificial neural networks consisting of two stages: encoding and decoding. A single-layer autoencoder (Fig. 3) is a kind of neural network consisting of only one hidden layer. Autoencoders aim at finding a common feature basis from the input data. They reduce dimensionality by setting the number of extracted features to be less than the number of inputs. Autoencoder models are usually trained by backpropagation in an unsupervised manner. The underlying optimization problem aims to minimize the distance between the reconstructed results and the original inputs.


Fig. 3. Autoencoder neural network architecture with one layer

The encoding function of an autoencoder (such as the one depicted in Fig. 3) maps the input data $s \in \mathbb{R}^d$ to a hidden representation $y \in \mathbb{R}^h = e(s) = g(w \cdot s + b)$ where $w \in \mathbb{R}^{d \times h}$ and $b \in \mathbb{R}^d$ are respectively the weight matrix and the bias vector and $g(\cdot)$ is a non linear activation function such as the sigmoid or hyperbolic tangent functions. The decoding function maps the hidden representation back to the original input space according to $\hat{s} = d(y) = g(w' \cdot y + b')$, $g(\cdot)$ being most of the time the same activation function.

The objective of the autoencoder model is to minimize the error of the reconstructed result:

$$(w, b, w', b') = \operatorname{argmin} \ell(s, d(e(s))) \tag{1}$$

where $\ell(u, v)$ is a loss function determined according to the input range, typically the mean squared error (MSE) loss:

$$\ell(u, v) = \frac{1}{n} \sum ||u_i - v_i||^2 \tag{2}$$

2) A new term of Regularisation: The method proposed in [2] is based on the analysis of the distribution of reconstruction errors $\rho_i = ||u_i - v_i||^2$. We found that once the autoencoder has converged, the reconstruction errors are distributed according to a (possibly exponential) rapidly decreasing density function with most samples centred around zero. The few samples with the highest reconstruction errors were associated with exceptional events (mostly weather related) whereas reconstruction errors located in the "belly" of the distribution happened to match traditional ATC deconfliction or sequencing operations.

A major challenge of our approach is the attempt to identify flows from trajectories entering and leaving a sector rather than from city-pair trajectories. Identified clusters may result from the aggregation of sparsely distributed trajectories. When samples are distributed in such a way that variation modes (or sub-clusters) of unbalanced weights emerge, these modes may subsist in the distribution of reconstruction errors in the form of a distribution with two "hills" (see Fig. 4). In order to limit a premature optimisation of the autoencoder which would learn to favour one mode over the other, we improved the training of our network with a regularisation term added to the loss $\ell(u, v)$.

A regularisation term is a penalty term added to the loss that is commonly used to prevent overfitting. Among neural networks, L1- (resp. L2-) regularisations penalize the loss after each iteration with the sum of the absolute (resp. squared) weights of the neural network. In our case, since we expect a distribution of reconstruction errors that would fit an exponential law, we propose a regularisation term based on a measure of distance between distributions.

After each iteration, we fit an exponential law to our distribution of reconstruction errors. The best fit to an exponential distribution can be written based on the mean of all (ρ_i) samples, which has already been computed in form of the MSE loss $\ell(u, v)$. Therefore, the best fit for the probability density function becomes:

$$f : x \mapsto \frac{1}{\ell(u, v)} \cdot e^{-\frac{x}{\ell(u, v)}} \tag{3}$$

Then, we compute the distance (Fig. 4) between the distribution of reconstruction errors and the fitted exponential probability density function. For each $t_j \in [0, \max(\rho_i)]$ equally sampled with $j \in [1, m]$, we evaluate the difference:

$$\delta_j = \left(\frac{1}{n} \sum_i \mathbb{1}_{[t_j, t_{j+1}]}(\rho_i) \right) - f(t_j) \tag{4}$$

Finally, we sum all the positive δ_j as a regularisation term to the original square loss. For our specific use case, we found $\lambda = 10^{-2}$ to be particularly efficient.

Fig. 4. Distribution of ρ_i (suggesting two modes of variations) and their fitted exponential law: the regularisation aims at minimising the coloured area.

$$\ell^\star(u,v) = \ell(u,v) + \lambda \sum_{j=1}^{m} \max(0, \delta_j) \qquad (5)$$

The differentiation of the regularisation term $\ell^\star(u,v)$, necessary to implement the gradient descent and backpropagation during the training period has been delegated to the autograd module of PyTorch. All terms presented in this subsection can be written with torch functions which provide all that is needed for backpropagation.

V. APPLICATION TO TRAJECTORY ANOMALY DETECTION

In order to identify the flows in the sector LFBBPT, we have applied the clustering method described in Section IV-A. This method is based on the DBSCAN algorithm, which requires two main parameters determining the size n and the density ε of clusters. For the first application of DBSCAN, we have set ε to 0.4 and n to 1% of the total number of trajectories. For the refinement of the clusters, DBSCAN has been executed with ε set to 0.5 and n set to 1% of the total number of trajectories in the cluster where it is applied. The minimum number of trajectories for a cluster to be formed has been established to 2% of the traffic in the sector.

The resulting cluster centroids representing the flows are displayed in Fig. 5. A total number of nine flows have been identified with a percentage of outliers reaching 26.4% of the traffic. We have checked how well the generated clustering centroids match the ATS Route Network (ARN) also published on eAIP. Some clusters fit well to sections of the published air routes, e.g. UN869 (cluster 1), UM728 (cluster 5) or UN460 (cluster 7) but a similar match is less evident with clusters 0, 3 and 4. Cluster 0 is an evolving flow with flights taking off from Paris area, so it seems reasonable to have it separated from clusters 3 and 4 which are both stable flows with a centroid at FL360. The reason for separating clusters 3 and 4 is less obvious, but certainly due to the fact that the exit points of these two clusters are separated by a less dense area, which can be observed in Fig. 1.

All trajectories within a cluster are then considered independently. We applied to each cluster the anomaly detection technique presented in Section IV-B1 with details in [2]. We


Fig. 5. The clustering method applied to our dataset trimmed to LFBBPT sector resulted in the nine following clusters. The altitude below the identifier of the cluster reflects the altitude profile of the centroid.

used a different autoencoder network architecture illustrated in Fig. 6, adapted to the resampling of our trajectories at 50 points per trajectories, working only with true track angles so as to focus on lateral resolutions of potential conflicts. The ADAM optimizer was iterated for 30,000 iterations.


Fig. 6. Neural network architecture with two layers of progressive compression (encoding) and two layers of decompression (decoding).

The MSE loss converged properly to each cluster, although the distribution of reconstruction errors of clusters with more sparsely distributed trajectories (e.g. cluster 4) lead to distribution profiles suggesting two modes, as reflected in Fig. 4. Using the regularisation term presented in Section IV-B2 lead to better results as Fig. 7 reflects: the same network has been trained twice on the same data, first with a regular MSE loss

$\ell(u,v)$, then with a regularised loss $\ell^{\star}(u,v)$; the MSE loss has been computed for both executions so as to provide a meaningful comparison.


Fig. 7. Convergence of the MSE loss is at first slower but soon becomes better with regularisation (cluster 4, $\lambda = 10^{-2}$): final value for the MSE loss after 30,000 iterations is $2.79 \cdot 10^{-4}$ with regularisation but $2.98 \cdot 10^{-4}$ without.

Eventually, even though the optimizer focused on optimising the network so as to minimize the regularised loss $\ell^{\star}(u,v)$, it found a better vanilla MSE loss $\ell(u,v)$. This result validates the idea behind regularisation which consists in penalizing our criterion hoping we can avoid overfitting and converge toward a more robust solution.

VI. ANALYSIS OF RESULTS

The output of our anomaly detection technique is a score (namely the reconstruction error) associated with each trajectory. The score is relevant with respect to the distribution of all scores inside the cluster. Results presented in [2] suggest that very high scores are to be related to exceptional events such as cumulonimbus impacting the area whereas samples located closer to the end of the distribution are often to be associated to ATC actions for deconfliction. This section analyses the result on the LBBPT sector and further validates this assumption.

We considered hereafter only a subset of our trajectories, namely trajectories with a reconstruction error higher than a given threshold. Since no reference catalogue of anomalous situations validated by experts is available, we have no choice but to arbitrarily set this value: we defined a threshold based on the fitted exponential distribution f defined in (3) and illustrated in Fig. 8 with the set of samples $\{x_i\}$ s.t.

$$f(x_i) \leq \frac{1}{5} \cdot f(0) \text{ i.e. } x_i \geq \log(5) \cdot \ell_{cluster}(u,v) \qquad (6)$$

A. Weather-related events

In order to validate our first assumptions about the most abnormal situations that were detected, we selected the top 10 trajectories with the highest reconstruction errors for each cluster, for a total of 90 trajectories. Table I shows that two days were particularly represented in that subset of trajectories. The last column (rank) reflects the position of the sample in the distribution: 1 stands for the highest reconstruction error, 2 for the second highest, etc.


Fig. 8. Selection of the sample trajectories with the higher reconstruction scores (example of cluster 4). For the purpose of the study, we chose a threshold of $1/5 = 20\%$ although this could be reconsidered in the future.

callsign	date and time of entry in LFBBPT	cluster	rank
TRA47R	2017-06-20 18:21:00Z	2	6
RYR9TG	2017-06-20 19:09:28Z	4	1
TOM84T	2017-06-20 19:14:51Z	4	3
RYR79EY	2017-06-20 18:35:46Z	4	7
SAA235	2017-06-20 19:07:15Z	5	9
DAH1007	2017-06-20 18:33:34Z	6	8
DLH68F	2017-06-20 18:25:17Z	7	8
VLG83TJ	2017-06-20 18:21:00Z	8	2
AFR88DM	2017-07-08 19:55:01Z	0	1
VLG8248	2017-07-08 20:09:53Z	0	6
AEA1008	2017-07-08 19:39:50Z	0	9
AAF221	2017-07-08 20:08:52Z	0	10
FIN611	2017-07-08 19:42:22Z	2	8

TABLE I
MOST SIGNIFICANT TRAJECTORIES/DAYS IN LFBBPT GROUPED BY DATE.

A first look at the METAR history in airfields located around the LFBBPT sector (LFLX, Châteauroux, to the North-West of LFBBPT and LFLC, Clermont-Ferrand in the Southern part of the sector) reflects locations of cumulonimbus (CB) and tower cumulus (TCU) consistent with the location of anomalous trajectories (Fig. 9): CB impacted the whole sector (hence clusters 2, 4, 5, 6, 7 and 8) on June 20th but only the Northern part of the sector (mostly clusters 0) on July 8th.

A more thorough assessment of the meteorological situation is available in Appendix. METAR history is printed for reference. Anomalous trajectories are also plotted together with the locations of cumulonimbus on June 20th, estimated from thermal IR data collected from the Spinning Enhanced Visible and Infrared Imager (SEVIRI) by the Meteosat Second Generation of Satellites.

B. Cross-analysis between flows

Isolating deconfliction ATC orders in regular traffic is a difficult task because most flights are executed without much deviation from their original intention. Reconstruction errors help isolate flights calling for further analysis. In an attempt to automatize the process, we computed the closest point of approach (CPA) for all pairs of trajectories which fly at the

Fig. 9. Situations with strong thunderstorm activities, on June 20th (left-hand side, reported in LFLX and LFLC) and July 8th (right-hand side, reported only in LFLX).

same moment in LFBBPT and which belong to our subset of trajectories defined in (6). For the CPA computation, we used the distance between two trajectories based on the cylindrical norm defined in [32]:

$$d_{CPA} = \min_{t} \left(\max \left(\frac{d_{lat}(t)}{5\,\text{nm}}, \frac{d_{vert}(t)}{1,000\,\text{ft}} \right) \right) \quad (7)$$

where d_{lat} is the distance between the two WGS84 coordinates and d_{vert} the difference of altitudes. 5 nm and 1000 ft are respectively the lateral and vertical separation minima required between aircraft flying within Reduced Vertical Separation Minima (RVSM) airspace [33]. Since we observe traffic which has allegedly been deconflicted by ATC, all pairs of trajectories should be separated by a distance $d_{CPA} \geq 1$. However, we assume that an action of deconfliction is likely to involve pairs of trajectories with a d_{CPA} relatively small.

In the following we focus on pairs of trajectories respecting the following conditions:

1) each trajectory has a reconstruction error higher than $\log(5) \cdot \ell_{cluster}(u, v)$ (see Fig. 8);
2) their $d_{CPA} \leq 2$, i.e. the lateral and vertical distance at the CPA should be smaller than 10 nm and 2000 ft respectively. In addition, we impose a constraint on the vertical distance to be smaller than 1500 ft in order to focus only on aircraft flying at adjacent flight levels.

Each trajectory being associated with a cluster, we build the density matrix as shown in Fig. 10 which reads as follows: the darker the color at position (i, j) with $i \geq j$, the more trajectories from cluster i and j are possibly subject to a deconfliction order from the ATC. A first consistency cross-check with the map on Fig. 5 seems convincing: cluster 0 interact with clusters 3 and 4 (mostly cluster 3); cluster 1 interact with clusters 2, 6 and 8, albeit less with cluster 5. Cluster 7 fly at a relatively lower altitude and only interact

with cluster 6 (trajectories climbing) but not with clusters 2 and 8 (constant altitudes).


Fig. 10. Density matrix (upper triangle only): the darker the color in (i, j), the more trajectories from clusters i and j are possibly subject to a deconfliction order from the ATC.

We focus in the following on specific situations in converging flows, then on pairs of trajectories in the same flow which may be impacted by the same factors.

C. Anomalies in pairs of trajectories from different flows

Fig. 11 reflects two situations involving cluster 1 and one of clusters 2, 6 and 8, containing mainly deconfliction situations for aircraft flying at the same level. Fig. 12 focuses on the interaction between cluster 0 (aircraft taking off from Paris area) integrating into en route traffic from clusters 3 and 4.


Fig. 11. Deconfliction actions between cluster 1 and cluster 6 (left, on January 3rd), resp. cluster 1 and cluster 2 (right, on March 4th).

The first situation of Fig. 11 involves TVF021Z (cluster 6) and TAP933A (cluster 1). The dashed line projects the situation as if no order had been given five minutes before the closest

point of approach with a possible conflict situation at the intersection of both routes. On the right hand side, ATC orders given to IBK7VY and IBK6113 seem to have anticipated the situation earlier with a probable deconfliction order given ten minutes before the closest point of approach. For such situations of converging routes at the same flight level, a future direction for improvement could be to automatically detect the level of anticipation of the deconfliction by looking backward from the CPA.

Situations involving cluster 0 are more representative of how to insert trajectories at the end of their climbing phase, at the moment they interact with en route flights. The first situation in Fig. 11 involves DAH1087 from cluster 0, taking off from Paris Orly and RYR24JE. The lateral separation between the two aircraft comes as close as 2 nautical miles before they start flying apart after 13:22 UTC. When DAH1087 reaches FL320, their separation seems comfortable (around 10 nm), but ATC probably gave a first clearance to FL320 before giving another clearance to climb to their final altitude of FL350, making sure the trajectories were properly separated. On the right hand side, IBE34PP, taking off from Paris Orly, and IBK6113 seem to have received proper deconfliction orders, yet IBE34PP is stopped at FL360 before being given an additional clearance to FL370.

On the left hand side, AEA1038 takes off from Paris Charles de Gaulle and RAM781S from Paris Orly. They both belong to cluster 0 and will probably be vectored on the same route. Having similar climb profiles (both aircraft are B738), special attention is paid to their lateral separation, probably leading to these peculiar trajectories.

On the right hand side, TAP817 flies Milan–Porto while CES709 flies Shanghai–Madrid. When both aircraft join route UN869, a special attention is paid to their separation (see the plateau at FL360) before they are probably separated by being placed on lateral offsets from UN869.


Fig. 13. Deconfliction situations between trajectories in the same clusters.

VII. CONCLUSIONS

In this paper, we presented a significant improvement to our previous contribution [2] based on the work in [1]. We developed an efficient methodology for detecting abnormal situations with a rule of thumb to distinguish weather related actions from deconfliction actions.

The approach applies an anomaly detection method to the results of a clustering algorithm. It yields convincing results, with the regularisation term presented in Section IV-B2 being a significant asset to address sparse clusters. The results of the method give a good grasp of the structure of traffic in LFBBPT sector when it is deployed by the ACC; a good taxonomy of problematic situations is covered by the results of this technique (thunderstorms, flows converging at the same flight level, one or several climbing trajectories to merge into en-route traffic, etc.) Such an analysis may give hints so as to find a better way to evaluate the workload associated to a sector.

Future work could include a robustness assessment so as to evaluate how stable the catalogue of resulting critical situations

Fig. 12. Deconfliction actions between cluster 0 and cluster 3 (left, on January 15th), resp. cluster 0 and cluster 4 (right, on March 31th).

D. Anomalies in pairs of trajectories from the same flows

The density matrix on Fig. 10 reflects a high number of potential deconfliction situations between pairs of trajectories from cluster 0 and from cluster 1. Fig. 13 looks into those situations.

is after several different trainings of the autoencoder, and after tuning the hyperparameters we introduced in the clustering and neural network training. Also, outlying trajectories from the clustering were left aside in this study although a thorough study of their interaction with the remaining traffic would probably be of great interest. Safety studies could also benefit from quantitative assessments resulting for our approach, possibly with a cross analyses with TCAS Traffic Advisories (TA) and Resolution Advisories (RA).

APPENDIX

A. METAR history

1) June 20th, 2017:

```
LFLX 202000Z AUTO 15002KT CAVOK 23/17 Q1015 TEMPO 18015G30KT
3000 -TSRA SCT040CB=
LFLX 201930Z AUTO 15003KT CAVOK 25/18 Q1015 TEMPO 18015G30KT
3000 -TSRA SCT040CB=
LFLX 201900Z AUTO 01003KT CAVOK 26/17 Q1015 TEMPO 18015G30KT
3000 -TSRA SCT040CB=
LFLX 201830Z AUTO 04005KT 350V080 9999 ///TCU 28/18 Q1015
TEMPO 18015G30KT 3000 -TSRA SCT040CB=
LFLX 201800Z AUTO 06008KT 020V110 9999 TS ///CB 31/17 Q1015
TEMPO 18015G30KT 3000 -TSRA=

LFLC 202000Z AUTO 28010KT 9999 RA ///TCU 25/16 Q1016 BECMG
NSW NSC=
LFLC 201930Z AUTO VRB03KT 9999 ///CB 29/15 Q1016 BECMG NSC=
LFLC 201900Z AUTO 35003KT 310V030 CAVOK 31/15 Q1016 NOSIG=
LFLC 201830Z AUTO 00000KT 9999 ///CB 32/14 Q1015 BECMG NSC=
LFLC 201800Z AUTO VRB02KT 9999 ///CB 31/16 Q1015 BECMG NSC=
```

2) July 8th, 2017:

```
LFLX 082100Z AUTO 10004KT 070V130 9999 -RA ///TCU 21/18
Q1015 BECMG NSC=
LFLX 082030Z AUTO 23013G25KT 9999 -RA VCTS ///CB 22/17 Q1016
BECMG 04008KT NSW NSC=
LFLX 082000Z AUTO 19004KT 110V210 9999 ///CB 22/17 Q1015
TEMPO 08015G25KT 2000 TSRA BECMG NSC=
LFLX 081930Z AUTO 07009KT 9999 ///CB 23/16 Q1014 TEMPO
08015G25KT 2000 TSRA BECMG NSC=
LFLX 081900Z AUTO 13013KT 090V160 9999 TS ///CB 23/16 Q1015
TEMPO 08015G25KT 2000 TSRA=

LFLC 082100Z AUTO 14003KT CAVOK 21/18 Q1016 TEMPO SCT040TCU=
LFLC 082030Z AUTO 26004KT CAVOK 19/17 Q1016 TEMPO SCT040TCU=
LFLC 082000Z AUTO 14003KT 120V180 CAVOK 20/18 Q1015 TEMPO
VRB15G25KT TSRA FEW040TCU=
LFLC 081930Z AUTO 31005KT 270V050 CAVOK 20/18 Q1015 TEMPO
VRB15G25KT TSRA FEW040TCU=
LFLC 081900Z AUTO 09010KT CAVOK 20/17 Q1017 TEMPO VRB15G25KT
2000 TSGR SCT040TCU=
```

B. Trajectories of aircraft flying to avoid cumulonimbus on June 20th 2017

Fig. 14 plots positions, trajectories for the past 15 minutes and comets for the next 5 minutes of aircraft flagged as anomalous by our method on June 20th 2017. The location of cumulonimbus is estimated from Thermal IR data from the Spinning Enhanced Visible and InfraRed Imager (SEVIRI), collected by the Meteosat Second Generation series of satellite.

The sector was implemented during the last hours of daylight, which may suggest that pilots avoid thunderstorms areas based on their visual perception in the early hours of operation, and relied on their on board weather radar in the later hours. Based on this kind of heatmap, future safety studies could assess how pilots manage to avoid thunderstorm areas based on the information of on board weather radars.

REFERENCES

[1] L. Basora, V. Courchelle, J. Bedouet, and T. Dubot, "Occupancy peak estimation from sector geometry and traffic flow data," in *Proc. of the 8th SESAR Innovation Days*, 2018.

[2] X. Olive, J. Grignard, T. Dubot, and J. Saint-Lot, "Detecting controllers actions in past Mode S data by autoencoder-based anomaly detection," in *Proc. of the 8th SESAR Innovation Days*, 2018.

[3] E. J. Garcia Gonzalez, "Development of a 3-dimensional mathematical collision risk model based on recorded aircraft trajectories to estimate the safety level in high density en-route airspaces," Ph.D. dissertation, Polytechnical University of Madrid, 2013.

[4] M. Gariel, A. N. Srivastava, and E. Feron, "Trajectory clustering and an application to airspace monitoring," *IEEE Transactions on Intelligent Transportation Systems*, vol. 12, no. 4, pp. 1511–1524, 2011.

[5] L. Basora, J. Morio, and C. Mailhot, "A trajectory clustering framework to analyse air traffic flows," in *7th SESAR Innovation Days*, 2017.

[6] X. Olive and J. Morio, "Trajectory clustering of air traffic flows around airports," *Aerospace Science and Technology*, vol. 84, Jan. 2019.

[7] S. Das, B. L. Matthews, A. N. Srivastava, and N. C. Oza, "Multiple kernel learning for heterogeneous anomaly detection: algorithm and aviation safety case study," in *Proc. of the 16th international conference on Knowledge discovery and data mining*, 2010, pp. 47–56.

[8] X. Olive and P. Bieber, "Quantitative assessments of runway excursion precursors using Mode S data," in *Proc. of the International Conference for Research in Air Transportation*, 2018.

[9] Y. Lü, Y. Duan, W. Kang, Z. Li, F.-Y. Wang *et al.*, "Traffic flow prediction with big data: A deep learning approach." *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 2, pp. 865–873, 2015.

[10] C. Di Ciccio, H. Van der Aa, C. Cabanillas, J. Mendling, and J. Prescher, "Detecting flight trajectory anomalies and predicting diversions in freight transportation," *Decision Support Systems*, vol. 88, pp. 1–17, 2016.

[11] Y. Liu, M. Hansen, D. J. Lovell, and M. O. Ball, "Predicting aircraft trajectory choice – a nominal route approach," in *Proc. of the International Conference for Research in Air Transportation*, 2018.

[12] M. Brittain and P. Wei, "Autonomous aircraft sequencing and separation with hierarchical deep reinforcement learning," in *Proc. of the International Conference for Research in Air Transportation*, 2018.

[13] C. Hurter, S. Puechmorel, F. Nicol, and A. Telea, "Functional Decomposition for Bundled Simplification of Trail Sets," *IEEE Transactions on Visualization and Computer Graphics*, vol. 24, no. 1, Jan. 2018.

[14] S. Puechmorel and F. Nicol, "Entropy minimizing curves with application to flight path design and clustering," *Entropy*, vol. 18, no. 9, 2016.

[15] A. T. Nguyen, "Identification of air traffic flow segments via incremental deterministic annealing clustering," Ph.D. dissertation, University of Maryland, 2012.

[16] G. Andrienko, N. Andrienko, G. Fuchs, and J. M. C. Garcia, "Clustering trajectories by relevant parts for air traffic analysis," *IEEE transactions on visualization and computer graphics*, vol. 24, no. 1, 2018.

[17] S. Rinzivillo, D. Pedreschi, M. Nanni, F. Giannotti, N. Andrienko, and G. Andrienko, "Visually driven analysis of movement data by progressive clustering," *Information Visualization*, vol. 7, no. 3-4, 2008.

[18] M. C. R. Murça, R. DeLaura, R. Hansman, R. Jordan, T. Reynolds, and H. Balakrishnan, "Trajectory clustering and classification for characterization of air traffic flows," *AIAA Aviation*, 2016.

[19] M. Ester, H.-P. Kriegel, J. Sander, X. Xu *et al.*, "A density-based algorithm for discovering clusters in large spatial databases with noise." in *Proc. of the 2nd International Conference on Knowledge Discovery and Data Mining*, vol. 96, no. 34, 1996.

[20] B. Matthews, A. N. Srivastava, J. Schade, D. R. Schleicher, K. Chan, R. Gutterud, and M. Kiniry, "Discovery of abnormal flight patterns in flight track data," in *2013 Aviation Technology, Integration, and Operations Conference*, 2013, p. 4386.

[21] B. Matthews, S. Das, K. Bhaduri, K. Das, R. Martin, and N. Oza, "Discovering anomalous aviation safety events using scalable data mining algorithms," *Journal of Aerospace Information Systems*, vol. 10, no. 10, pp. 467–475, 2013.

[22] B. Matthews, D. Nielsen, J. Schade, K. Chan, and M. Kiniry, "Automated discovery of flight track anomalies," in *Proc. of the 33rd Digital Avionics Systems Conference*, 2014.

[23] L. Li, M. Gariel, R. J. Hansman, and R. Palacios, "Anomaly detection in onboard-recorded flight data using cluster analysis," in *Proc. of the 30th Digital Avionics Systems Conference*, 2011.

[24] L. Li, S. Das, R. John Hansman, R. Palacios, and A. N. Srivastava, "Analysis of flight data using clustering techniques for detecting abnormal operations," *Journal of Aerospace information systems*, vol. 12, no. 9, 2015.

[25] A. Nanduri and L. Sherry, "Anomaly detection in aircraft data using recurrent neural networks (RNN)," in *Proc. of the Integrated Communications Navigation and Surveillance conference*, 2016.

[26] I. Melnyk, A. Banerjee, B. Matthews, and N. Oza, "Semi-Markov switching vector autoregressive model-based anomaly detection in aviation systems," in *Proc. of the 22nd International Conference on Knowledge Discovery and Data Mining*, 2016.

[27] I. Melnyk, B. Matthews, H. Valizadegan, A. Banerjee, and N. Oza, "Vector autoregressive model-based anomaly detection in aviation systems," *Journal of Aerospace Information Systems*, 2016.

[28] S. Das, B. L. Matthews, and R. Lawrence, "Fleet level anomaly detection of aviation safety data," in *Conference on Prognostics and Health Management (PHM)*, 2011.

[29] W.-H. Lee, J. Ortiz, B. Ko, and R. Lee, "Time series segmentation through automatic feature learning," *arXiv:1801.05394*, 2018.

[30] T. Dubot, "Predicting sector configuration transitions with autoencoder-based anomaly detection," in *Proc. of the International Conference for Research in Air Transportation*, 2018.

[31] M. Schäfer, M. Strohmeier, V. Lenders, I. Martinovic, and M. Wilhelm, "Bringing up OpenSky: A large-scale ADS-B sensor network for research," in *Proc. of the 13th international symposium on Information processing in sensor networks*, 2014, pp. 83–94.

[32] C. A. Munoz and A. J. Narkawicz, "Time of closest approach in three-dimensional airspace," NASA, Tech. Rep., 2010.

[33] A. ICAO, "unified framework for collision risk modelling in support of the manual on airspace planning methodology with further applications, circ 319-an/181 ed," *International Civil Avaition Organization, Montreal, Canada*, 2008.

ACKNOWLEDGMENT

AUTHORS' BIOGRAPHIES

Xavier Olive graduated from the French Institute for Aeronautics and Space (ISAE–SUPAERO) in Toulouse in 2006 and received the Ph.D. degree in informatics from the University of Kyoto in 2011. He is a full-time researcher with ONERA DTIS and specializes in operations research, data analytics and machine learning applied to flight data analytics, air traffic management and operations.

Luis Basora graduated in 2005 with a double master's degree in computer science and ATM from both the Polytechnic University of Catalonia (UPC) and the French Civil Aviation School (ENAC). He has worked since 2007 as an ATM research engineer in ONERA DTIS and he has been involved in SESAR since 2010 in collaboration with the French ANSP DSNA. Currently, his main areas of interest in research include data analytics and machine learning applied to air traffic data.
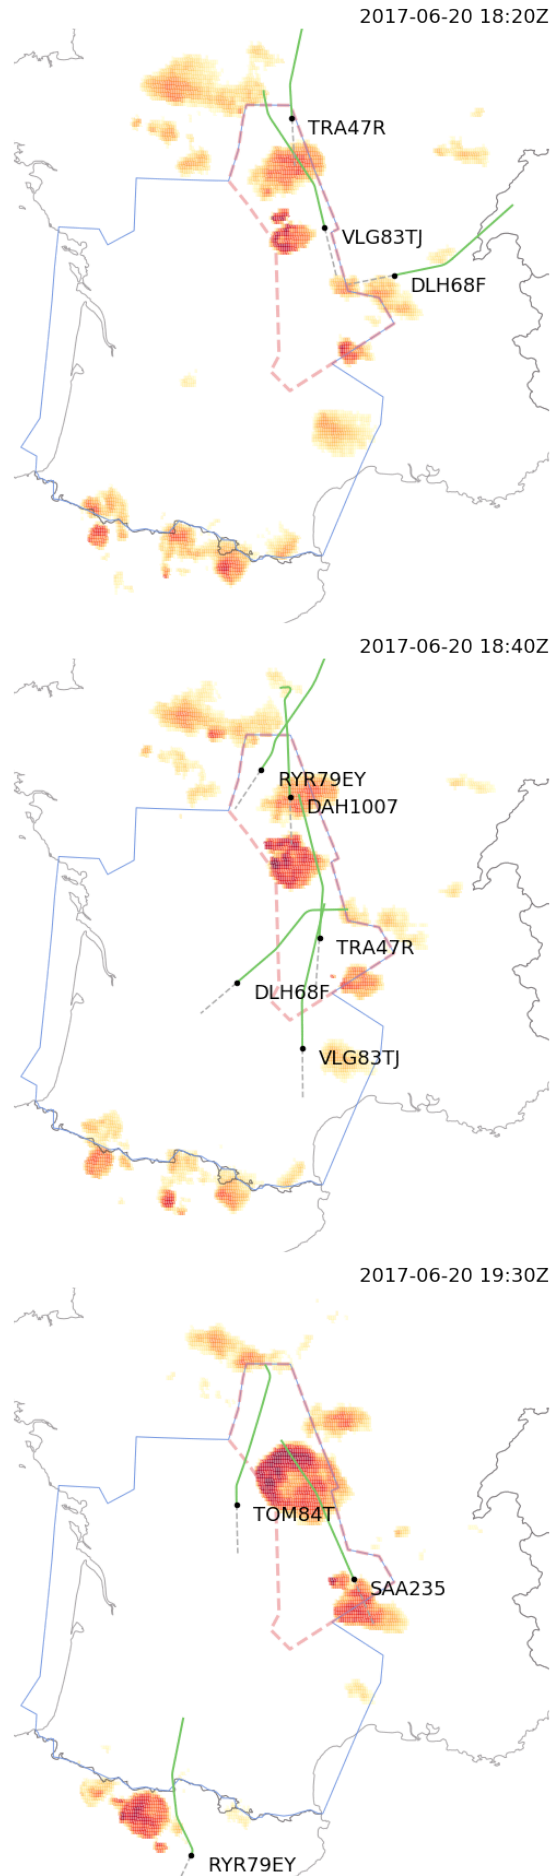
Fig. 14. Trajectories of aircraft flying to avoid cumulonimbus