Robust Aircraft Trajectory Planning under Uncertain Convective Environments with Optimal Control and Rapidly Developing Thunderstorms

Daniel González-Arribas, Manuel Soler, Manuel Sanjurjo-Rivo, Maryam Kamgarpour, Juan Simarro[a,b,c,d,e]

[a]*Department of Bioengineering and Aerospace Engineering, Universidad Carlos III de Madrid, Leganés, Spain. e-mail: daniel.gonzalez.arribas@ing.uc3m.es (corresponding author). Avenida de la Universidad, 30, Leganés (28911 Madrid), Spain*
[b]*Department of Bioengineering and Aerospace Engineering, Universidad Carlos III de Madrid, Leganés, Spain. e-mail: masolera@ing.uc3m.es*
[c]*Department of Bioengineering and Aerospace Engineering, Universidad Carlos III de Madrid, Leganés, Spain. e-mail: msanjurj@ing.uc3m.es*
[d]*Automatic Control Lab, ETH Zurich, Zurich, Switzerland. e-mail: mkamgar@control.ee.ethz.ch*
[e]*Agencia Estatal de Meteorología (AEMET), Valencia, Spain. e-mail: jsimarrog@aemet.es*

Abstract

Convective weather, and thunderstorm development in particular, represents a major source of disruption, delays and safety hazards in the Air Traffic Management system. Thunderstorms are challenging to forecast and evolve on relatively rapid timescales; therefore, aircraft trajectory planning tools need to consider the uncertainty in the forecasted evolution of these convective phenomena. In this work, we use data from a satellite-based product, Rapidly Developing Thunderstorms, to estimate a model of the uncertain evolution of thunderstorms. We then introduce a methodology based on numerical optimal control to generate avoidance trajectories under uncertain convective weather evolution. We design a randomized procedure to initialize the optimal control problem, explore the different resulting local optima, and identify the best trajectory. Finally, we demonstrate the proposed methodology on a realistic test scenario, employing actual forecast data and an aircraft performance model.

Keywords: optimal control, robust planning, convective weather, aircraft trajectory planning

1. Introduction

Convective weather, in particular thunderstorm development, creates major challenges for the Air Traffic Management (ATM) system. Flying through a thunderstorm entails multiple risks to the aircraft, such as strong turbulence, wind shear, downbursts, icing, lightning or hail. Consequently, aircraft should avoid them, and pilots are thus instructed to deviate from their flight plans to avoid strong convective activity. These deviations, when aggregated over multiple flights, amount to significant air traffic disruptions. Such disruptions must then be managed by Air Traffic Control (ATC) and Air Traffic Flow Management (ATFM) authorities, often by holding, delaying or cancelling other flights.

Just to provide a quantification of its impact, authors in [1] estimate the cost of ATFM delay in 100 Euros per minute. Looking into data, it should be remarked that weather phenomena added 0.61 minutes of primary delay[1] to each flight taking place in the European Civil Aviation Conference (ECAC) airspace in 2017, resulting in around 6.5 million minutes of primary delay[2], amounting to 17% of all primary delay not due to airlines. For the same year in the US, the data from the Federal Aviation Administration's OPSNET portal[3] attributes around 60% of the delays greater than 15 minutes per flight (which represent a total of 22.6 million minutes) to weather. Convective weather is one of the leading causes, with almost half of weather related delays in the busy summer months being caused by it.

A key reason behind the disruption caused by thunderstorms is the difficulty of forecasting their birth and evolution with precision at flight planning timescales (in this paper, 1-3 hours before departure). Therefore, taking them into account represents a challenge for flight dispatchers, air traffic controllers

---

[1] "Primary delay" refers to delay that is not reactionary, i.e., it is not attributable to late arrivals of aircraft, crew or connecting flights because of delays earlier in the day.

[2] CODA Digest 2017: http://www.eurocontrol.int/sites/default/files/publication/files/coda-digest-annual-2017.pdf

[3] https://aspm.faa.gov/opsnet/sys/Delays.asp

and flow management authorities. While some meteorological conditions are required for thunderstorm formation and can be forecasted in advance [2], the specific location and timing of its initiation is harder to identify. The characteristic sizes and lifespans of thunderstorms are small compared with the spatiotemporal resolution of medium-range Numerical Weather Prediction (NWP) models, and the chaoticity of the atmospheric dynamics compounds the challenge.

As a consequence, both thunderstorm weather forecasting and avoidance take place at shorter timescales. Prediction is usually performed in the form of deterministic nowcasts[4]. Instead of NWP based methods, thunderstorm nowcasting has been traditionally based on extrapolation of radar echoes [3], which has been improved by the usage of satellite data and Doppler radars. This is due to the fact that extrapolation is more effective than NWP at short timescales (1 to 2 hours). Nonetheless, the quality of extrapolation nowcasts degrades rapidly as the forecasting horizon increases (and particularly so for smaller convective cells). One extrapolation based nowcasting system of particular interest for aviation is the Corridor Integrated Weather System (CIWS)[4], which covers the Conterminous United States (CONUS) and is in use by the Federal Aviation Administration (FAA).

More recently, extrapolation based nowcasting methods have started to include NWP forecasts in order to improve their accuracy beyond the 1 to 2-hour mark, where the statistical skill of extrapolation techniques in nowcasting drops below that of NWP forecasting due to the characteristic timescales of convection initiation, growth and decay [5, 6]. These combined systems blend the extrapolation and a high-resolution NWP forecast whose initial conditions have been improved by data assimilation. By placing higher weight in the extrapolation for

---

[4]Within a meteorological context, "nowcasting" refers to forecasting for short ranges (6 hours at most, according to the World Meteorological Organization (WMO)) at the mesoscale (weather phenomena of sizes between a few kilometers and a few hundred kilometers), and relying more on extrapolation of current sensor data (radar echoes, mainly) than on NWP techniques.

short horizons and the NWP forecast for longer horizons, it is possible to combine the strengths of both approaches. One such combined system, designed for aviation purposes, is CoSPA[7]. It integrates CIWS data with forecasts from the National Oceanic and Atmospheric Administration (NOAA) and the National Center for Atmospheric Research (NCAR). Additionally, the trend towards increasing computational power has increased the interest in higher resolution, convection-permitting NWP forecasts in the last years. In [8], the AROME ensemble forecast from Météo-France is employed in conjunction with statistical post-processing techniques in order to forecast reflectivity (a variable that is highly correlated to convective weather hazards) in a probabilistic fashion. Nevertheless, the statistical performance of such blended systems is still limited and research is ongoing at meteorological centres around the world to improve in this aspect.

A different trend is the study of thunderstorm forecasting with data-driven methods with or without extrapolation, eschewing the use of NWP modelling. In [9], the authors conducted a characterization of the uncertainty in the movement of thunderstorms, as detected by a radar-based system, for usage in aircraft routing problems. Approaches to convection forecasting based on machine learning are also starting to be employed; for example, classifiers with Support Vector Machines (SVM) [10] or predictors with neural networks[11].

Finally, an alternative aviation-oriented approach focuses on analyzing and forecasting "weather avoidance fields", i.e., regions of space that pilots will try to avoid. This can be performed by using statistical analysis to identify relationships between processed convection-related meteorological variables and weather-caused deviations of aircraft from the planned route. An important system of this kind is the Convective Weather Avoidance Model (CWAM) [12, 13, 14], covering the CONUS. The CWAM assigns a probability of deviation to each point in space according to spatially filtered values of the Vertically Integrated Liquid (VIL) metric and the difference between the flight altitude and the echo top of the storm. The work in [15] develops probabilistic deviation thresholds based on the NCWF-6 convective forecasts. Thus, CWAM has been

4

extensively studied and employed in US-based ATM research, but no model of this kind has been developed to the same extent for other regions of the world.

A different approach is taken in [16]; instead of employing flown trajectories as data source, a forecast-based scenario is presented to pilots, who are then asked to design avoidance trajectories in order to characterize pilot behaviour. It was found that pilots are willing to take more safety margins when it is easy to do so and fly closer to the storms when avoidance is difficult. The authors conclude that there might be a trade-off between efficiency and comfort that could be derived empirically.

All in all, a first problem of interest is to develop a data-driven probabilistic model for thunderstorm phenomena based on existing (deterministic) products in Europe. A second problem of interest is to incorporate such probabilistic forecast in trajectory planning tools in order to design deviation routes that avoid the such potentially unsafe zones.

The flight planning literature describes multiple kinds of algorithms that have been employed for this purpose. The simplest of them are based on geometric procedures [17, 18, 19]. This class of methods do not generally consider thuderstorm evolution, uncertainty, or trajectory optimality; however, their simplicity results in fast computational times (thus being compatible with real-time usage) and easy implementation and integration with other tools and algorithms.

A second class of methods is based on graph-search algorithms such as Dijkstra's shortest path algorithm, A^* or D^*. These methods feature good computational properties too, but it is harder to model time-varying costs (such as fuel burn) and aircraft dynamics. In [20], the authors combine such a procedure with a multi-objective genetic algorithm in order to produce a set of proposed reroutes around convective weather, as analyzed by CWAM. In [21], CWAM weather forecasts were employed in combination with Dijkstra's algorithm to minimize a combination of fuel burn and expected cost of deviation due to weather. In [22], a 3-D field D^* (which avoids the optimality loss arising from the restriction to discrete legs and, thus, discrete headings) is applied in a receding horizon fashion in order to minimize a combination of flight time and

5

accumulated probability of bad weather conditions along the route.

Alternatively, methods that consider aircraft dynamics explicitly can be employed. In [23], a receding horizon optimization scheme is employed in order to compute avoidance trajectories. In [24], uncertainty in the dynamics and wind is modelled through a stochastic differential equation and the resulting optimal control problem is solved by discretizing the state space to obtain a Markov chain, which is then optimized through Jacobi iteration. A different approach involves formulating a stochastic reach-avoid problem [25, 26, 27], which is then solved through dynamic programming techniques. These works rely on state space discretization and thus are vulnerable to the "curse of dimensionality", i.e., the exponential scaling of the required computational resources with the dimension of the problem.

None of the methods studied in the literature is yet able to generate realistic, model-based trajectories that take into account both aircraft dynamics and uncertain thunderstorm evolution with real-time or near-real-time computational times. Therefore, in this paper we aim to develop a methodology for aircraft trajectory planning in the presence of convective thunderstorms whose evolution is considered uncertain. We will consider single-aircraft trajectory segments with a duration of 15 to 90 minutes, corresponding to the convective weather encounter. We address the problem of replanning the trajectory while the aircraft is already flying, but around 5 to 45 minutes before the aircraft would encounter severe adverse weather.

To that end, this study makes the following contributions. First, we develop and fit a model of the uncertain motion of the convective weather cells on top of a deterministic nowcasting system based on extrapolation, Rapidly Developing Thunderstorms (RDT), which covers Europe and neighbouring regions and can be deployed in other regions with Meteosat coverage. We then make use of the produced probabilistic nowcast in order to formulate an optimal control problem whose solution is an optimized aircraft trajectory that minimizes a combination of flight cost (in terms of both time and fuel) and the probabilistic exposition to convective weather conditions. We also introduce a randomized

6

initialization scheme for the optimization algorithm that allows us to explore different local optima of the problem, with the goals of identifying the best one and also providing alternative trajectory options. Finally, we demonstrate the capabilities of the proposed methodology on a realistic scenario.

We employ numerical control techniques that allow us to generate optimal trajectories while considering nonlinear aircraft and thunderstorm dynamics at a moderate computational cost. In addition, numerical optimal control is powerful and versatile enough that potential future extensions (such as varying flight altitudes) can be incorporated with manageable increments of the computational cost.

We work within a free-routing framework, which is more appropriate for a thunderstorm encounter than a structured airspace approach where the aircraft trajectory is restricted to fly through certain airways and waypoints. Additionally, we work on a robust control paradigm ("design a single trajectory and control history that is robust in the presence of uncertainty") instead of a stochastic control paradigm ("design a control policy that generates different trajectories depending on the realization of the uncertainty"). The latter concept leads to better performance, as the control actions are not taken a priori and are thus chosen with more information; however, we consider robust control to be more appropriate for ATM purposes. The main reason is that fixed and predictable trajectory plans make it easier for air traffic controllers to coordinate multiple aircraft trajectories, which is particularly important in this setting as convective weather episodes lead to higher controller workload.

The proposed solution could find applications in both onboard and ground-based decision-support tools (DST). At the airborne side, pilots could employ such a DST to quickly obtain and evaluate flight path alternatives; in the ground, it could help air traffic controllers to assign time-based metering slots to aircraft flying through weather impacted regions and/or group aircraft trajectories into weather-avoiding flows that are easier to manage. Finally, it could be employed as a basis for trajectory synchronization and negotiation between air traffic control and the aircraft.

7

This work is structured as follows. In Section 2, we describe the characterization of convection, from the data source to the probabilistic model. The optimal control formulation of the problem, as well as its initialization and solution are discussed in Section 3. We show the results of using our approach on a realistic scenario in Section 4 and draw conclusions on Section 5.

2. Meteorological modelling

Convection occurs when the stratification of the atmosphere breaks down due to the presence of convective instability, with the formation of tower clouds that easily reach the tropopause. It is observed over a wide range of spatial and temporal scales [28]:

- At annual time periods and planetary-length scales, convection is modulated by meteorological events governed by orbital and oceanic forcings, such as the Intertropical Convergence Zone (ITCZ), the monsoonal circulations and the El Niño Southern Oscillation (ENSO).

- At the synoptic scale and middle latitudes, Rossby waves produce low and high pressure systems with length scales from 1000 to 6000 km and durations on the order of one or two weeks. While high pressure systems tend produce a stable stratification of the atmosphere that inhibits convection, low pressure systems tend to feature convection along cold, warm and occluded fronts and instability lines [29].

- At the mesoscale, we can find convective phenomena with characteristic lengths of 10 to 1000 km, such as isolated thunderstorms, squall lines, convective complexes and tropical cyclones, that can be embedded in bigger structures. [30].

While planetary and synoptic scale phenomena can be forecasted with a global Ensemble Prediction System (EPS) [31, 32], these forecast systems lack the spatiotemporal resolution to resolve mesoscale phenomena. Limited area

8

EPS products (which nowadays have resolutions near the kilometer) can be employed at the mesoscale [33, 34], but they are not yet able to forecast individual cells.

Mesoscale phenomena (in particular, thuderstroms) are the main focus of the trajectory planning problems that we consider in this study. Therefore, as discussed in Section 1, it is useful to employ a nowcasting system based on extrapolation instead of an NWP product. We employ the RDT system[5] for this purpose, which we describe in Section 2.1.

2.1. Data

RDT is one of the products of the Satellite Application Facility on support to Nowcasting and Very Short-Range Forecasting (NWCSAF) consortium, participated by several European National Weather Services. It uses imagery collected from the SEVIRI instrument installed aboard the geostationary Meteosat Second Generation (MSG) satellites, with a horizontal resolution of 3 km, in order to characterize convective systems in an area around Europe every 15 minutes.
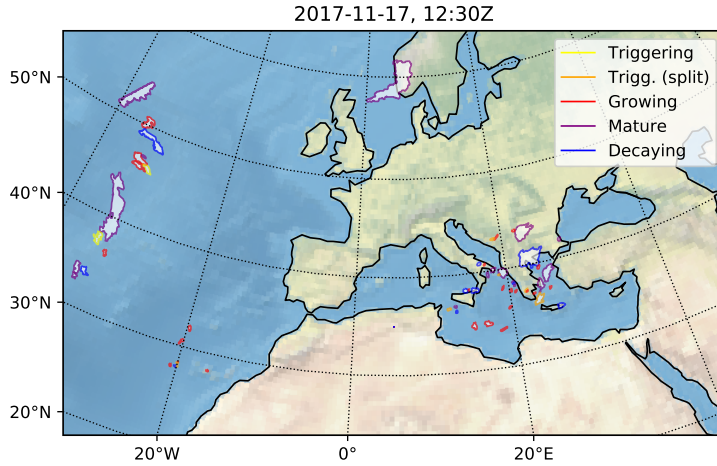

Figure 1: Convective systems over Europe as identified by the RDT product.

The RDT system identifies convective cells and describes them as polygons on a latitude-longitude map in an object-oriented approach, while also determining certain attributes of the cells (see Figure 1). The algorithm at the core of RDT operates in three steps:

- In first place, an *identification* of cells corresponding to cloud systems is performed using the infrared brightness temperature channel from the SEVIRI data. The algorithm identifies the cloud towers from the local minima of the infrared brightness temperature and then analyzes the local temperature pattern in the vicinity of the cloud tower to determine an adaptive temperature threshold that is employed to determine the extent of the cloud system.

- In second place, successive analyses are compared with a *tracking* algorithm that employs overlap and correlation methods in order to match previously determined cells with the newly identified objects and store their evolution and trajectories; this information is employed to generate new estimates of the speed and direction of the cells. Cloud splitting and merging events are also handled at this stage.

- Finally, a *discrimination* method is employed to distinguish convective objects from other cloud cells, which are much more numerous. It relies on data from multiple infrared and vapour channels and uses both spatial characteristic and temporal variations and trends. There parameters are stored in a learning database, which is used to train a statistical model that employs lightning occurrence data from the Météorage and European Cooperation of Lightning Decision (EUCLID) networks as ground truth.

The RDT output contains the list of identified convective objects, along with characteristics such as the perimeter, the speed and direction of its motion, or the cloud top pressure. One computed parameter of particular interest for our analysis is the *phase* of the convective cell, which classifies thunderstorm

10

systems into five categories: triggering, triggering from split, growing, mature and decaying.

These parameters (in particular, the perimeter, speed and direction) can be employed to extrapolate the position of the storm in a deterministic manner (indeed, RDT output includes extrapolated nowcasts generated in this fashion). However, this extrapolation is deterministic and doesn't take into account the error inherent to the nowcast. Therefore, we proceed to build a model of the uncertainty in Section 2.2.

2.2. Probabilistic modelling

We will now seek to obtain a function of the form $p_{t_a}(\mathbf{r}, t)$ that represents the probability that the point $\mathbf{r} = (\phi, \lambda) \in \mathbb{R}^2$ (where ϕ and λ denote geodetic latitude and longitude in degrees) lies inside a convective storm at time t, as estimated using the forecasts available up to time t_a.

Let t_a be the time at which we build the probabilistic forecast (not the deterministic RDT forecast). We define the reference time $t_{ref} <= t_a$ as the time at which the latest RDT forecast before t_a was released.[6] See Figure 2 for reference.

As input data, we will make use of the latest N_{steps} forecasts that are available before t_a. We will denote the corresponding forecasts times as follows (see Figure 2):

$$t^f_{-i} = t_{ref} - i \cdot 15 \text{ min}, \, i \in \mathcal{I} := \{0, 1, \ldots, N_{steps} - 1\}$$

For example, if $N_{steps} = 3$ (as in our implementation), we will employ the forecasts made at t_{ref}, $t_{ref} - 15$ min and $t_{ref} - 30$ min.

For each of these forecasts, the projected position of each storm can be extrapolated to any posterior time t in a linear and deterministic manner by

---

[6]While each analysis contains forecasts for multiple discrete time horizons (15, 30, 45, and 60 minutes), we can extrapolate them to any instant between them by making use of the speed and direction information.

Figure 2: Construction of the time-lagged ensemble.

employing the velocity information provided by the forecast and the temporal distance. In order to perform the extrapolation, we add the quantity $\mathbf{v} \cdot |t - t^f_{-i}|$ to each point in the polygon that defines the contour of the storm, where $\mathbf{v}$ is the estimated velocity of the storm in latitude-longitude coordinates and can be expressed as:

$$\mathbf{v} = \frac{180}{\pi} R_e^{-1} v_s \begin{bmatrix} \cos \chi_s \\ (\cos \phi)^{-1} \sin \chi_s \end{bmatrix} \tag{1}$$

where R_e is the mean radius of the Earth, v_s is the velocity of the storm as estimated by the RDT forecast and χ_s is the estimated direction of movement of the storm (again provided by the RDT extrapolation).

Let $\mathcal{J} = \{$Triggering, Triggering from split, Growing, Mature, Decaying$\}$ be the set of possible phases of the storm. For a given reference time, we now compute the indicator functions $I_{-i,j}(\mathbf{r}, t)$ with $i \in \mathcal{I}, j \in \mathcal{J}$. These functions take the value 1 if, according to the forecast made at t^f_{-i} and extrapolated to

12

time t in the aforementioned fashion, the point $\mathbf{r}$ lies inside a storm with phase j; they take the value 0 otherwise.

We will now proceed to incorporate stochasticity into the movement of the convective cell in order to represent the forecast uncertainty.[7] We model this evolution as a Brownian motion with drift. This process is described by the following stochastic differential equation (see [35, Chapter 5]):

$$d\mathbf{c}_t = \mathbf{v}dt + \Sigma_w d\mathbf{W} \tag{2}$$

In Equation (2), $\mathbf{c}_t$ denotes the position of a reference point in the storm, such as the center; $\mathbf{v}$ denotes the speed of the storm. Σ_w is a matrix describing the magnitude and shape of the random perturbations; we model them as isotropic, i.e., $\Sigma_w = \sigma_w \mathbb{I}_{2\times 2}$ with σ_w denoting the amplitude of the perturbations. Finally, $\mathbf{W}$ denotes a standard two-dimensional Wiener process. The Wiener process (also known as "Brownian motion") represents the simplest stochastic process; it is characterized by having independent increments, distributed as Gaussian variables with variance equal to the time differential (i.e. with the information available at time t, the value of the Wiener process at $t + \Delta t$ is a normal variable with variance Δt). It can, therefore, be viewed as the "integral in time" of white noise.

We define the stochastic displacement $\mathbf{D}$ as the difference between the stochastic position $\mathbf{c}_t$ and the deterministic forecast $\mathbf{c}_{t^f_{-i}} + \mathbf{v}(t - t^f_{-i})$ (see Figure 3). For a fixed t, this displacement is a Gaussian random variable with mean 0 and standard deviation $\sigma_w \sqrt{t - t^f_{-i}}$; its probability density function is, therefore:

$$p_{\mathbf{D}}^t(\mathbf{y}) = \frac{1}{\sqrt{2\pi}\sigma_w |t - t^f_{-i}|} \exp\left(-\frac{||\mathbf{y}||^2}{2\sigma_w^2 |t - t^f_{-i}|}\right)$$

Because the position of the storm is now a stochastic process, the indicator functions $I_{-i,j}(\mathbf{r}, t)$ become stochastic processes too. We can, however, compute

---

[7]In this work, we will not consider changes in size or shape

Figure 3: Example realization of Equation (2).

their expected value $\bar{I}_{-i,j}(\mathbf{r}, t)$ with a straightforward convolution operation, as shown by equation (3). See Figure 4 for a schematic illustration of the convolution process.

$$\bar{I}_{-i,j}(\mathbf{r}, t) = \mathbb{E}[I_{-i,j}(\mathbf{r}, t)] = \int_{\mathbb{R}^2} p_{\mathbf{D}}^t(\mathbf{y}) I_{-i,j}(\mathbf{r} - \mathbf{y}, t)) d\mathbf{y} = p_{\mathbf{D}}^t * I_{-i,j}(\cdot, t) \quad (3)$$

Figure 4: Schematic representation of the deterministic and stochastic (averaged) indicator functions in one dimension at different, successive times.

Performing this approximation for multiple storms in the same manner requires the assumption that the stochastic displacement is common to all storms with the same phase j; in other words, the movement of the storms is assumed to be perfectly correlated and linearity allows us to take the convolution after aggregation. While it would be desirable to model the correlation of the movements of the storms, we consider this approximation good enough for our purposes, as it approximates a "correct" $\bar{I}_{-i,j}$ by an upper bound.[8]

At this stage, we have $|\mathcal{I}| \times |\mathcal{J}|$ smoothed, extrapolated indicator functions. In order to integrate the information from all of them in a data-driven fashion, we will use them as predictors in a statistical model. We select a logistical

---

[8]Assume that there are two storms, and the events that the point $\mathbf{r}$ is in each one of them at time t are denoted by A and B. Then

$$P(A \cup B) = P(A) + P(B) - P(A \cap B) = \bar{I}_{-i,j}(\mathbf{r}, t) - P(A \cap B).$$

Therefore, if $P(A \cap B)$ (which represents the probability of the point being inside both extrapolated storms) is small enough, the approximation is close. For more storms, the argument is similar, as the difference is again only intersection terms.

regression model, as it is a parsimonius specification that models the output variable as the probability that the event of interest takes place (in this case, the existence of convective conditions in a future step). Therefore, by comparing the predicted indicators with the actual outcomes (according to the analysis in the next time steps), a logistical regression procedure can be employed to predict convective conditions in a probabilistic manner.

In a Generalized Linear Model (GLM) such as linear regression, an outcome variable y (that, in a binary setting like ours, takes values in $\{0, 1\}$) is modelled as a random variable whose expected value depends on the values of a vector $\mathbf{z}$ of predictors or "features":

$$\ell(\mathbb{E}[y]) = \mathbf{w}^T \mathbf{z} \tag{4}$$

where ℓ represents a *link function* while $\mathbf{w}$ is a vector of coefficients whose values are determined, in GLM routines, by maximizing their log-likelihood on a training set. The quantity $\mathbf{w}^T \mathbf{z}$ is called the *linear predictor*.[9] In the case of logistic regression, the link function is the *logit* function:

$$\mathrm{logit}(p) = \log\left(\frac{p}{1-p}\right) \tag{5}$$

In our case, the predictor variables will be the extrapolated indicator functions as well as a constant offset w_0, so the linear predictor is given by:

$$\mathrm{LP}(\mathbf{r}, t) = w_0 + \sum_{\substack{i \in \mathcal{I} \\ j \in \mathcal{J}}} w_{i,j} \bar{I}_{-i,j}(\mathbf{r}, t) \tag{6}$$

and the desired probability function can be written as:

$$p_{t_{ref}}(\mathbf{r}, t) = \frac{1}{1 + \exp(-\mathrm{LP}(\mathbf{r}, t))} \tag{7}$$

Section 2.3 describes the process of fitting this model.

---

[9]Note that, just like in the linear regression case, nonlinear predictors can be included as long as they are included in additive fashion.

The complete model is defined by the parameters N_{steps} and the coefficients σ_w, w_0 and $w_{i,j}$. We choose $N_{steps} = 3$ after our exploratory analysis concluded that the coefficients after $i = 2$ were too close to 0 to be relevant and the statistical skill of the model was not appreciably deteriorated by the exclusion of the $i > 2$ forecasts; in other words, the information contained in the forecasts before the third to last seems to be redundant for prediction purposes. In order to determine the remaining parameters, we will make use of three procedures:

- An "inner loop" that computes the values of w_0 and $w_{i,j}$ for a given value of σ_w, as well as a measure of the statistical skill of the model.

- An "outer loop" that calls the inner loop for different values of σ_w and then selects the value σ_w^* that leads to the best statistical skill.

- A "bootstrap-like" method that runs the inner loop with $\sigma_w = \sigma_w^*$ and different, randomly selected slices of the dataset in order to obtain more robust estimates of the w_0 and $w_{i,j}$ coefficients as well as to estimate their volatility.

We start by describing the inner loop procedure. For a given value of σ_w, we can compute the values of the coefficients w_0 and $w_{i,j}$ with the following method:

1. Select a number of dates and times t_a in random fashion from a training dataset and load the corresponding RDT analyses, as well as the N_{steps} preceding forecasts.

2. For each analysis, compute the "areas of interest" (AoI), defined as rectangular regions containing any detected or forecasted storms while adding proportional margins (see Figure 2.3). This step reduces the subsequent computational requirements of the process.

3. At each point $\mathbf{r}$ in a rectangular grid inside an AoI, compute the values of $\bar{I}_{-i,j}(\mathbf{r}, t_a)$ and the actual value $I_a(\mathbf{r}, t_a)$ (i.e., 1 if the point lies inside a storm according to the analysis and 0 otherwise). Note that the

17

extrapolated and smoothed indicator $\bar{I}_{-i,j}(\mathbf{r}, t_a)$ depends on the value of σ_w.

4. Launch a logistic regression routine employing the computed values of $\bar{I}_{-i,j}$ as predictors and the values of the points as data samples, returning the values of the w_0 and $w_{i,j}$ coefficients.

The performance of the regression model can be evaluated with several metrics. Suppose that we now consider the test set (composed by data that have not been employed in training) and predict that every data point where the
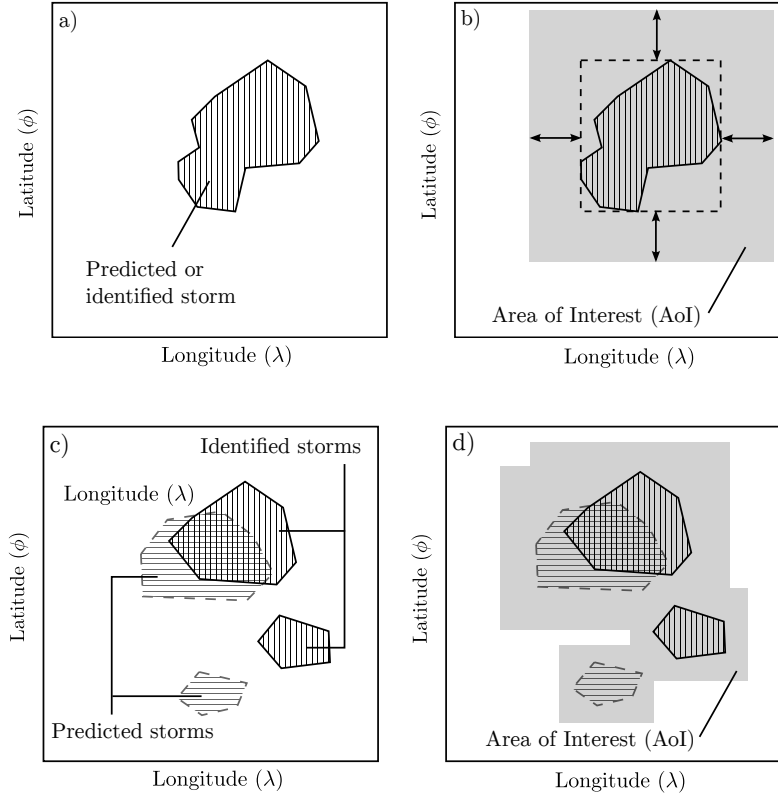

Figure 5: Area of Interest determination. For a given storm (as in Figure 2.3.a), the corresponding AoI is the rectangle containing it, with proportional margins (Figure 2.3.b). For each analysis, comprised of multiple identified storms plus forecasted storms from earlier analysis (Figure 2.3.c), the AoIs from each storm are merged.

18

forecasted probability is above a threshold to have convective conditions and every other data point to have clear air conditions. If we plot the true positive rate against the false positive rate as the threshold changes from 0 to 1, we obtain the Receiver Operating Characteristic (ROC) curve. The integral of the area below the ROC curve is called the Area Under Curve (AUC) metric; AUC values of around 0.5 indicate that the model is no more skilled that random classification on average, while values closer to 1 indicate that the skill of the model is high.[10]

We can now employ this inner loop process in order to choose the value of σ_w in the outer loop and the values of w_0 and $w_{i,j}$ in the bootstrap-like procedure. Both methods are illustrated in Figure 6.

For the outer loop, we can employ the AUC criterion to choose the "best" value of σ_w. We compute the logistic regression on a training set for different values of the σ_w parameter (using the same randomly chosen data). Then, we compute the AUC score on a test set. Finally, we select the value of σ_w that lies on the maximum of the σ_w - AUC curve. This process is performed twice: first, with a logarithmic sweep of σ_w in order to identify the scale of interest (upper plot in Figure 7); then, with a linear sweep and polynomial approximation in order to identify the optimum with more precision (though we will note that somewhat lower or higher values lead to very similar AUC scores). With our dataset, comprising RDT analysis and forecasts from the 20th to the 29th of March, 2017, the chosen value is $\sigma_w = 0.73 \ \mathrm{km} \cdot \mathrm{s}^{-0.5}$.

---

[10]Values closer to 0 indicate "reverse skill", i.e., the model is worse than chance at classifying outcomes so the reverse prediction has skill.

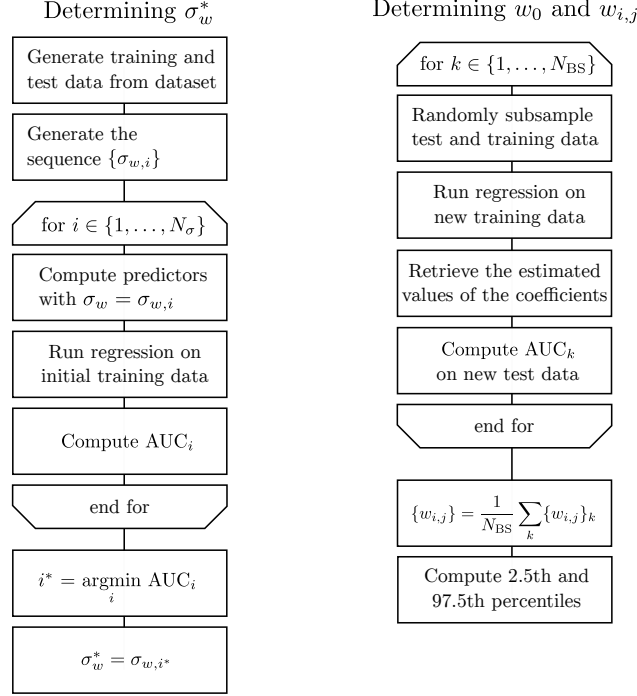Determining σ_w^*

Generate training and
test data from dataset

Generate the
sequence $\{\sigma_{w,i}\}$

for $i \in \{1, \ldots, N_\sigma\}$

Compute predictors
with $\sigma_w = \sigma_{w,i}$

Run regression on
initial training data

Compute AUC_i

end for

$i^* = \underset{i}{\arg\min}\ \text{AUC}_i$

$\sigma_w^* = \sigma_{w,i^*}$

Determining w_0 and $w_{i,j}$

for $k \in \{1, \ldots, N_{\text{BS}}\}$

Randomly subsample
test and training data

Run regression on
new training data

Retrieve the estimated
values of the coefficients

Compute AUC_k
on new test data

end for

$\{w_{i,j}\} = \frac{1}{N_{\text{BS}}} \sum_k \{w_{i,j}\}_k$

Compute 2.5th and
97.5th percentiles

Figure 6: Schematic representation of the deterministic and stochastic (averaged) indicator functions in one dimension at different, successive times. N_{BS} denotes the number of bootstrap analyses.

Finally, with σ_w^* fixed, we can now estimate the mean values and confidence intervals of the coefficients $\{w_{i,j}\}$ with the bootstrap-like procedure. We randomly select 250 dates and times from the available data set and subsample 25% of the resulting areas of interest at random; these settings allow us to hold the resulting sample in-memory. Then, we obtain the coefficient values by regression and evaluate the performance by computing the AUC score on a test dataset composed by a random 25% of a random 150 dates and times from the remaining files. We repeat this process 1000 times and take the average, as well as the 2.5th and the 97.5th percentiles which allow us to define the 95% confidence intervals.

Figure 7: Values of the AUC (Area Under Curve) metric. The dashed line represents the polynomial approximation employed to determine the value of σ_w that optimizes the AUC.

The computed values of the coefficients $\{w_{i,j}\}$ are presented in Table 1 and Figure 8. It can be observed that the step of the forecast is more relevant in the prediction than the analyzed storm phase, with the latest forecast being 2 to 4 times more influential than the previous one, while the earliest forecast is statistically significant only partially. Additionally, we can see that the estimated values of the coefficients for storms that are triggering are more volatile; we attribute this outcome to the fact that they are less numerous within the data set, so the coefficient estimates have higher variance.

	0 min	-15 min	-30 min
Triggering	5.712 (4.955, 6.793)	2.556 (1.474, 3.544)	0.052 (-0.868, 1.229)
Tr. (split)	6.453 (6.114, 6.813)	1.756 (1.199, 2.342)	-0.054 (-0.676, 0.575)
Growing	5.647 (5.397, 5.886)	1.239 (0.875, 1.613)	-0.082 (-0.430, 0.282)
Mature	5.765 (5.518, 6.026)	1.746 (1.380, 2.097)	0.699 (0.346, 1.030)
Decaying	5.700 (5.493, 5.920)	1.555 (1.231, 1.873)	0.581 (0.248, 0.914)

$$w_0 \quad | \quad \text{-4.155 (-4.186, -4.122)}$$

Table 1: Coefficient values w_0 and $\{w_{i,j}\}$ and 95% confidence intervals.
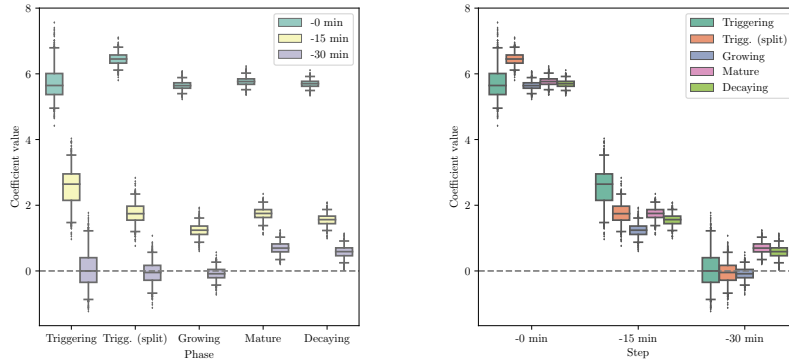
Figure 8: Coefficient estimates w_0 and $\{w_{i,j}\}$. Each boxplot represents the distribution among different bootstrap-like fits, with the box representing the 25th, 50th and 75th percentiles and the whiskers representing the 10th and 90th percentiles.

The computed AUC values range from 0.949 to 0.958. This seems like an excellent value, but we will note that our dataset features a large amount of points where no storm is extrapolated and no storm occurs; this inflates the AUC score by providing a lot of "easy predictions". In practice, our model will often miss newly created storms that are hard to detect using extrapolation procedures, unless they are close enough to other storms that the probability dispersion effect assigns some non-zero probability.

Figure 9 compares the deterministic extrapolation of an RDT nowcast to the probability field generated by the fitted model. It can be observed how the probability field diffuses as the forecast horizon increases.
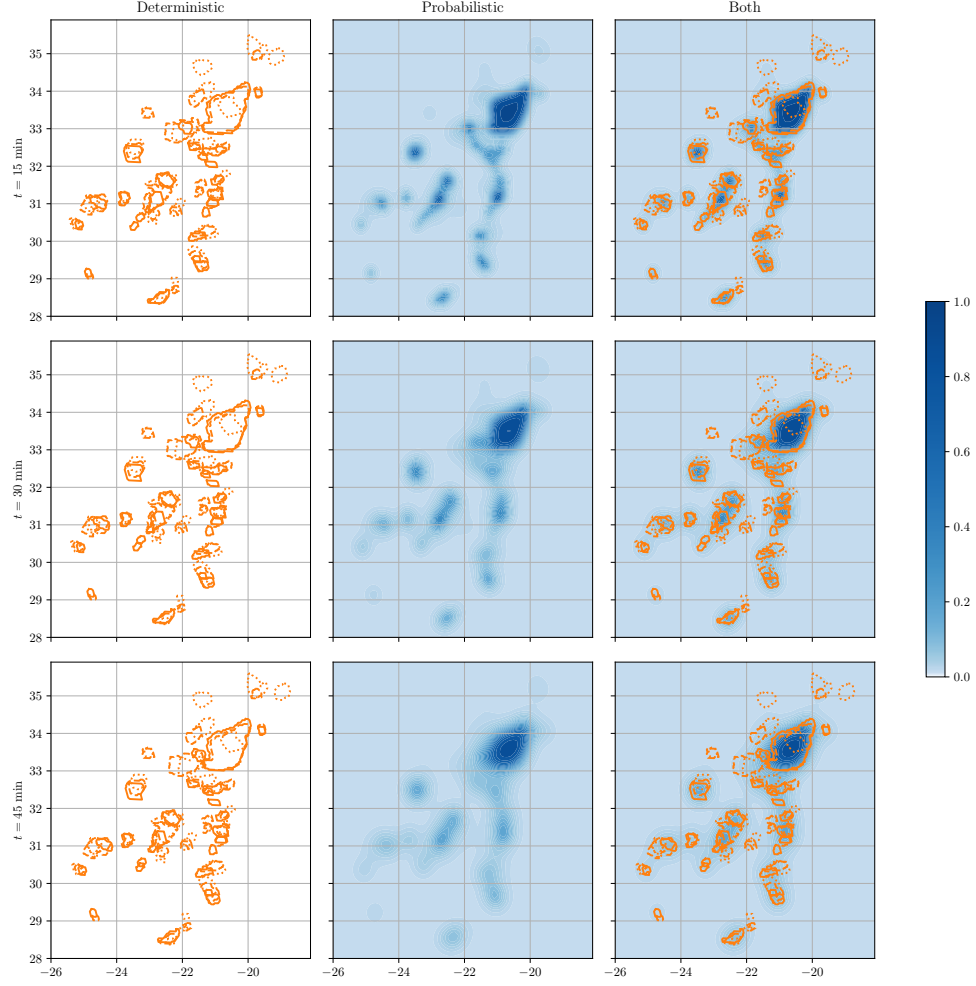

Figure 9: Comparison between the deterministic RDT extrapolation and the probabilistic model at different forecast horizons. The solid, dashed, and dotted lines represent the extrapolation from the forecast at t^f_{-0}, t^f_{-1}, and t^f_{-2}, respectively. Color bar denotes $p_t(\mathbf{r}, t)$.

We now turn our attention to the issue of how to employ these generated probabilistic forecasts for the purpose of flight planning under uncertainty.

3. Optimization Methodology

We use a numerical optimal control framework, as it represents a versatile approach that not only allows us to model nonlinear dynamics and nonconvex objective functions (features that have made it popular in aerospace engineering [36, 37]), but it is also flexible enough to allow for future extensions of the method.

3.1. Optimal Control

Consider a controlled, constrained dynamical system, composed by:

- A temporal domain $\mathcal{T} = [t_0, t_f] \subset \mathbb{R}$

- A vector of state variables $\mathbf{x}(t) : \mathcal{T} \to \mathbb{R}^{n_x}$

- A vector of control variables $\mathbf{u}(t) : \mathcal{T} \to \mathbb{R}^{n_u}$

- The differential equation $\mathbf{f} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathcal{T} \to \mathbb{R}^{n_x}$ that describes the evolution of the system:

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t) \tag{8}$$

- Equality constraints $\mathbf{h} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathcal{T} \to \mathbb{R}^{n_h}$ (with the equality sign applying in element-wise fashion):

$$\mathbf{h}(\mathbf{x}(t), \mathbf{u}(t), t) = 0 \tag{9}$$

- Inequality constraints $\mathbf{g} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathcal{T} \to \mathbb{R}^{n_g}$ (with the inequality signs applying in element-wise fashion):

$$\mathbf{g}_L \leq \mathbf{g}(\mathbf{x}(t), \mathbf{u}(t), t) \leq \mathbf{g}_U \tag{10}$$

In order to define an optimal control problem, we need to define a *cost functional* that describes the objective that must be pursued, as well as a set of *boundary conditions*. The boundary conditions can be defined by a function $\Psi : \mathbb{R}^{n_x} \times \mathbb{R} \times \mathbb{R}^{n_x} \times \mathbb{R} \to \mathbb{R}^{n_b}$ and an allowable set $\Omega_b \subset \mathbb{R}^{n_b}$ such that:

$$\Psi(\mathbf{x}(t_0), t_0, \mathbf{x}(t_f), t_f) \in \Omega_b \tag{11}$$

Finally, the cost functional is commonly written in the general "Bolza form":

$$J(\mathbf{x}, \mathbf{u}, t_0, t_f) = \Phi(\mathbf{x}(t_0), t_0, \mathbf{x}(t_f), t_f) + \int_{t_0}^{t_f} \mathcal{L}(\mathbf{x}(t), \mathbf{u}(t), t) dt \tag{12}$$

where $\Phi : \Omega_x \times \mathbb{R} \times \Omega_x \times \mathbb{R}$ is often called the "Mayer term" or "terminal cost" (as it usually depends only on the final time and state) and the integral term is called the "Lagrangian term" or "running cost", with $\mathcal{L} : \Omega_x \times \Omega_u \times \mathcal{T} \to \mathbb{R}$. The optimal control problem can then be defined as:

$$\min_{t_0, t_f, \mathbf{x}, \mathbf{u}} J(\mathbf{x}, \mathbf{u}, t_0, t_f) \tag{13}$$

such that Equations and inequalities (8) – (11) are fulfilled.

The solution of an optimal control problem is characterized by either the necessary conditions for optimality (the Pontryagin Minimum Principle, which transform the optimal control problem into a two-point boundary value problem) or the sufficient conditions (described by the Hamilton-Jacobi-Bellman partial differential equation). Most optimal control problems of practical interest cannot be solved analytically, so numerical methods are employed instead. We choose to employ direct transcription methods [38], as they handle constraints more naturally than indirect methods and do not suffer from the "curse of dimensionality" like dynamic programming.

In direct collocation methods [39], the state and control trajectories are discretized in time and the differential equations are turned into *defect constraints* that depend on the values of the states and controls at the discretization nodes, which become decision variables. The cost functional is approximated by a quadrature rule, and the constraints are enforced at the discretization nodes. In this form, the optimal control problem can be turned into an Nonlinear Programming (NLP) problem, which can then be solved with standard nonlinear optimization algorithms such as SNOPT [40] or IPOPT [41], which make use of gradient-based methods for solving the optimization problem.

3.2. Dynamical Model

We assume that the uncertainty in the movement of the storm is substantially greater than the uncertainty in aircraft dynamics; therefore, we model the motion of the aircraft as a deterministic system. We assume the aircraft moves in a 2D plane (i.e. at a constant flight level) with variable airspeed and non-instantaneous turning dynamics. We consider an ellipsoidal Earth as in the World Geodetic Survey 1984 reference system (WGS84) model, with radii of curvature of ellipsoid meridian and prime vertical denoted by R_M and R_N respectively. We take wind from a forecast and set remaining atmospheric parameters according to the International Standard Atmosphere (ISA) model. The equations that describe the dynamics are:

$$\frac{d}{dt}\begin{bmatrix} \phi \\ \lambda \\ v \\ \chi \\ m \end{bmatrix} = \begin{bmatrix} (R_N + h)^{-1}(v\cos(\chi) + w_x(\phi, \lambda, t)) \\ (R_M + h)^{-1}\cos^{-1}(\phi)(v\sin(\chi) + w_y(\phi, \lambda, t)) \\ (\mathrm{T}(C_T) - \mathrm{D}(C_L, v))/m \\ (g/v)\tan\mu \\ -\eta(v)\mathrm{T}(C_T, v) \end{bmatrix}, \tag{14}$$

$$\mathrm{L}(C_L, v)\cos\mu = mg, \tag{15}$$

where ϕ is the latitude, λ is the longitude, v is the true airspeed, m is the mass, h is the geodetic altitude, χ is the heading, μ is the bank angle of the aircraft, w_x and w_y are the zonal and meridional components of wind,[11] T is the thrust force, D is the drag force and η is the thrust-specific fuel consumption. Both η and D are modelled according to the BADA 4 aircraft performance model [42], and $\mathrm{L}(C_L, v) = \frac{1}{2}\rho v^2 S C_L$ is the lift force, where S is the wet wing surface. The control vector is composed by the thrust coefficient C_T, the lift coefficient C_L, and the bank angle μ. Figure 10 illustrates the relationship between the speeds and angles.

---

[11] Contrary to the usual definition, we take w_y to be in a South to North direction, so that the signs match.

Figure 10: Relationship between airspeed, groundspeed, wind, heading and course.

In addition, the following flight envelope constraints apply:

$$|\mu| \leq 45°, \tag{16}$$

$$C_L \leq C_{Lmax} \tag{17}$$

$$M(v) \leq M_{MO} \tag{18}$$

$$C_{T_{min}} \leq C_T \leq C_{T_{max}} \tag{19}$$

$$v \leq v_{max} \tag{20}$$

where C_{Lmax} is the maximum lift coefficient, M is the Mach number (the ratio between airspeed and the speed of sound, see Equation 21), M_{MO} is the divergence Mach number, $C_{T_{min}}$ and $C_{T_{max}}$ are the thrust limits and v_{max} is a tighter airspeed limitation that we introduce to account for the potential speed limits when flying through turbulent air near convective zones. The speed of sound, a, depends on the temperature, so the Mach number can be defined as:

$$M = v/a(T) \tag{21}$$

27

3.3. Trajectory Optimization

We formulate the trajectory planning problem as an optimal control problem. Its objective function is formulated as follows:

$$J = \left[-m(t_f) + \text{CI} \cdot t_f + r_p \int_{t_0}^{t_f} p(\mathbf{r}(t), t) dt \right], \tag{22}$$

where CI is the "Cost Index" parameter, representing the preference for reduced flight time instead of fuel burn, and r_p is the "risk penalty", which we define as the preference for reduced exposition to regions of high storm probability instead of the other objectives. A higher r_p setting generates a trajectory that is less likely to be exposed to convective activity.

The optimal control problem is formulated as:

$$\min J \tag{23}$$

subject to Equations (14) to (20) as well as the boundary conditions:

$$(\phi, \lambda, v, \chi, m)(t_0) = (\phi_0, \lambda_0, v_0, \chi_0, m_0) \tag{24}$$

and

$$(\phi, \lambda, v)(t_f) = (\phi_f, \lambda_f, v_f) \tag{25}$$

If a required time of arrival (RTA) to the end waypoint is specified, we set $t_f = t_{RTA}$ instead of leaving it free. Naturally, if an RTA is set in this fashion, the CI setting becomes irrelevant.

As we mentioned before, we work with direct transcription methods that transform the optimal control problem into a nonlinear optimization problem. Therefore, a user-provided initial guess for the decision variables (the values of the trajectory state and control variables at every node) is required. We describe the generation of this initial guess in the next subsection.

Note that the presented optimal control scheme requires a probability field of function of the form $p(\mathbf{r}(t), t) \in [0, 1]$ (in either analytical or gridded representation), such as the one we generate in Section 2.2. Therefore, a different weather

forecast or product could be employed as input, as long as it is provided in the required shape and represents the same concept (the probability of an aircraft at $\mathbf{r}$ being within a convective cell at time t) or a similar one.

3.4. Initialization

The selection of the initial guess is important for two reasons. First, the quality of the initial guess is, in general, an important factor in the performance of NLP solvers. The number of iterations, the computational time (particularly relevant for our purposes) and the likelihood that the algorithm converges to an optimal solution are all dependant on the initial guess.

In second place, we expect the optimization problem that we formulate to have multiple local minima; therefore, the initial guess will determine which of these multiple solutions will be found by the algorithm. We expect different local minima because of the non-convex nature of the problem, which is derived to a high extent from the existence of multiple potential avoidance routes that are separated (in trajectory space) by "worse" routes that cross the obstacles, which implies the existence of multiple basins of attraction around the locally optimal routes.

For these reasons, we have designed a randomized heuristic initialization procedure that produces different initial guess trajectories with the goal of exploring the solution space and find the different local minima. Generating multiple solutions also has operational benefits, as pilots and controllers can then choose one of them according to a different criterion (for example, facilitating deconfliction or sequencing); additionally, it is possible for them to employ their trained intuition and experience to judge the feasibility and complexity of the trajectories, or any factor not included in the model, to take the best decision in practice.

The generated trajectories will be obtained assuming no wind and constant airspeed v_{IG}. Under these assumptions, any trajectory departing from $\mathbf{r}_0$ can be represented by the arrival time t_f and the heading history $\chi(t)$. If a random arrival time t_f and a random heading history $\chi(t)$ are generated, integrating the dynamics leads to a random trajectory. However, this randomly-generated tra-

29

jectory will arrive at a random point, and we are only interested in trajectories that arrive at $\mathbf{r}_f$. Therefore, we will generate a "normalized" heading history $\hat{\chi} : [0, 1] \to \mathbb{R}$ such that

$$\chi(t) = \chi_0 + \hat{\chi}\left(\frac{t}{t_f - t_0}\right), \tag{26}$$

where χ_0 and t_f are adjustable constants that allow us to generate a valid heading history from any continuous normalized heading history $\hat{\chi}$ through a straightforward rotation (adjustment of χ_0) and stretch (adjustment of t_f).

In order to generate $\hat{\chi}$, we consider the following basis expansion:

$$\hat{\chi}(\tau) = \sum_{k=1}^{n_L} a_k Z_k P_k(\tau) \tag{27}$$

where n_L represents the degree of the expansion (we use 4), $\{Z_k\}$ is a sequence of independent standard normal variables, $\{a_k\}$ represents a weighting sequence (in our implementation, $a_k = (3/4)^{1-k}$) and $\{P_k\}$ denotes the shifted Legendre polynomials (the Legendre polynomials on $[0, 1]$). Note that the zeroth term is omitted, as it corresponds to a constant offset and χ_0 already fulfills that role. Table 3.4 lists the shifted Legendre polynomials up to degree 5.

k	$P_k(\tau)$
0	1
1	$2\tau - 1$
2	$6\tau^2 - 6\tau + 1$
3	$20\tau^3 - 30\tau^2 + 12\tau - 1$
4	$70\tau^4 - 140\tau^3 + 90\tau^2 - 20\tau + 1$
5	$252\tau^5 - 630\tau^4 + 560\tau^3 - 210\tau^2 + 30\tau - 1$

Table 2: Legendre polynomials on $[0, 1]$

By sampling $\{Z_k\}$ in random fashion, we can generate different functions $\hat{\chi}(t)$, which determine different heading histories $\chi(t)$. We generate the rest of the state and control trajectory by integrating Equation (14), thus completing

30

the initial guess. Finally, by solving the problem with different starting points obtained in this fashion and collecting the corresponding solutions, we have a higher chance of obtaining a global optimum instead of a single local optimum.

We note that the relationship between a given starting point and the optimal solution found by the solver is not fixed, but may depend on the choice of NLP solver and its settings.

3.5. Computational setup

We employ the Python-based CasADi [43] library for NLP modelling. The probability field is computed in a $0.1° \times 0.1°$ grid, sampled every 5 minutes, and then interpolated through 3-dimensional B-splines. We employ the interior-point NLP solver IPOPT [41] running with the MA27 sparse symmetric linear solver from the HSL Mathematical Software Library [44] and initial barrier parameter $\mu = 10^{-3.8}$. The computations are performed in a workstation equipped with an Intel Xeon E3-1240 v5 CPU running at 3.5 GHz.

We employ a trapezoidal transcription scheme [39, Chapter 4] with piecewise-constant controls as direct collocation method, with defect constraints of the form

$$\mathbf{x}_{n+1} - \mathbf{x}_n = \frac{h}{2} \left(\mathbf{f}(\mathbf{x}_{n+1}, \mathbf{u}_n, t_n) + \mathbf{f}(\mathbf{x}_n, \mathbf{u}_n, t_{n+1}) \right), \ \forall n \in \{0, \ldots, N_{\text{nodes}} - 1\},$$

with the subscript n denoting the value of the directized variables at the n-th node or interval and $h = \dfrac{t_f - t_0}{N_{\text{nodes}} - 1} \equiv t_{n+1} - t_n$ denoting the node spacing. By employing an homogeneous node spanning h, this transcription scheme samples the probability field in a homogeneous fashion, unlike higher-order pseudospectral methods [45]; furthermore, the potential existence of constrained arcs (such as legs at maximum speed) does not allow us to assume that pseudospectral methods will provide spectral accuracy unless coupled with complex adaptive methods. In any case, the choice of discretization method is not critical for our work.

4. Case study

4.1. Scenario definition

Our test scenario is based on a storm group that is detected by the RDT algorithm on November 16th, 2017. The storm group is evolving towards North-East. The scenario starts at 6:00Z and ends at 07:30Z. A twinjet narrow-body airliner modeled according to the BADA 4 specification flies at FL330 from $(\phi_0, \lambda_0) = (34°, -24°)$ to $(\phi_f, \lambda_f) = (28°, -19°)$, therefore crossing the group of storms. The remaining initial conditions are given by:

$$\chi_0 = 115° \tag{28}$$

$$m_0 = 71495 \text{ Kg} \tag{29}$$

$$v_0 = 220 \text{ m/s} \tag{30}$$

We select $r_p = 10$ Kg/s, a required time of arrival at the end waypoint of 07:12Z, a v_{max} of 230 m/s, and a final airspeed of 200 m/s. We solve the problem with the initialization procedure described in Section 3.4, starting from 3400 randomly chosen starting points. We will discuss the best solution found in Section 4.2 and the other local minima afterwards, in Section 4.3.
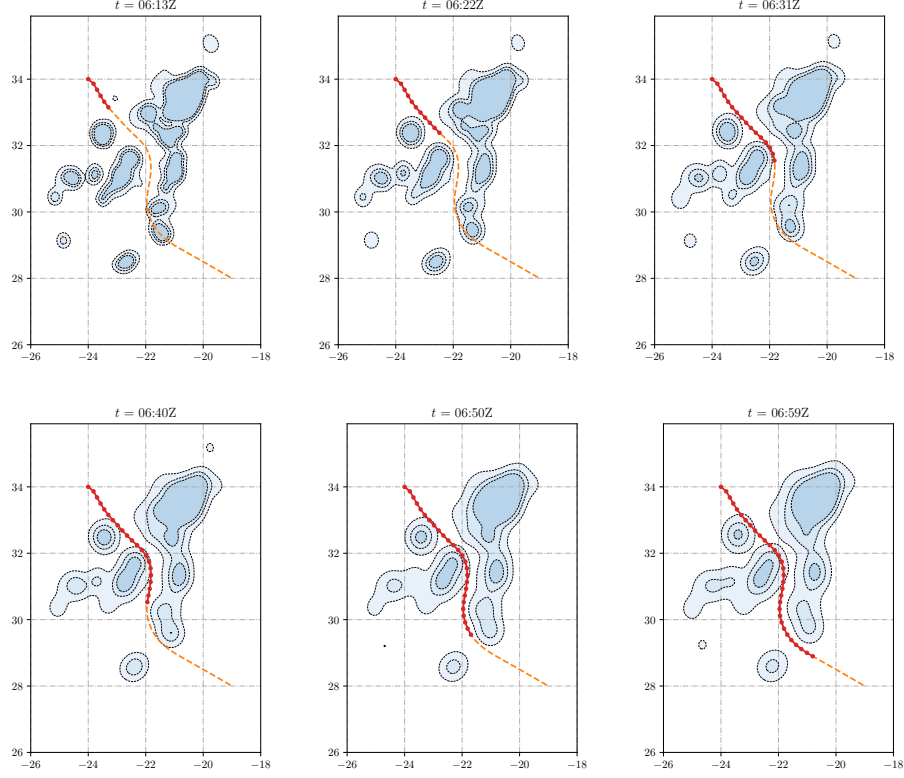
Figure 11: Computed trajectory and storm probability $p(\mathbf{r}, t)$ at different points in the trajectory. The solid line with dots represents the flown part of the trajectory while the dashed one represents the remaining part. The three contour levels represents the 2.5%, 5% and 10% probability levels.

Figure 11 illustrates the evolving probability field $p(\mathbf{r}, t)$ at different instants in time, as well as the planned trajectory of the aircraft, passing through a corridor between the storms. It can be observed that the planned trajectory takes the future potential evolution of the storms into account, and not just the present location of the storms.
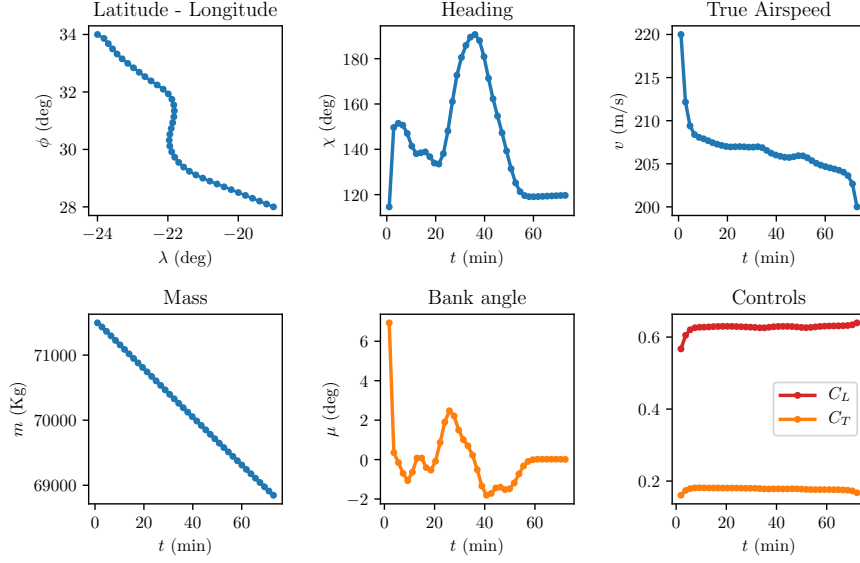
Figure 12: State space and control trajectories for the global solution.

State and control trajectories are illustrated in Figure 12. Since the resulting airspeed profile is relatively smooth, the longitudinal controls (C_T and C_L) are stable, while the lateral-directional control (the bank angle, μ, with C_L playing a compensation role) changes in order to steer the aircraft. Beyond the initial and final transitions towards the initial and final conditions, the speed slowly decreases as the mass decreases. This could be attributed to multiple factors: it might be optimal to go faster at the beginning in order to avoid the storms incoming from starboard and to go slower at the final leg in order to wait for the storms at the port side to go away. Alternatively, the TAS profile might be similar to the one in a simple, one-dimensional cruise problem [46]; in this kind of setting, the optimal solution requires that the airspeed slowly goes down as the mass of the aircraft decreases.

4.3. Local optima

Figure 13: Trajectories from different randomized starting points. For the most likely trajectories with close-to-optimal objective function values, we annotate the difference in cost with respect to the global optimum and the percentage of the time that the randomized initial guess produces the trajectory. The storm probability field is represented at 06:30Z.

We illustrate all the locally optimal paths found by the method in Figure 13. We highlight the five solutions that are present in more than 1% of the runs and have objetive values closer than 100 Kg to the global optimum.

Figure 14: TAS profiles of the randomly initialized trajectories

It can be observed that about two thirds of the starting points lead to the most common trajectory, which is very close in cost (12 Kg) to the global optimum. Another 4.4% of the starting points lead to the global optimum, as discussed in Section 4.2. Together, the five highlighted trajectories account for nearly 85% of the initialization points; the remaining 15% of the starting points lead to trajectories that have higher costs or are found infrequently. Some of these solutions try to route around the whole storm group, but the RTA requirement prevents them from fully avoiding the bigger storm clusters.

These results confirm our hypothesis about the multiplicity of local optima formulated in Section 3.4. Therefore, the usage of a randomized initial guess procedure is justified.

Figure 14 displays the airspeed profiles of these trajectories. The most common solution features the slowest airspeed profile and, thus, the lowest expected fuel burn (as long as the trajectory does not have to reroute). It does so, however, at the expense of increasing the potential exposition to convective hazards when compared to the global optimum (see Figure 15(a) and Figure 15(b)).

36

Another feature present in the airspeed profiles shown in Figure 14 is the appearance of peaks of increased airspeed, which correspond to the crossing of zones of higher storm probability (compare with Figure 13 and Figure 15(a)). This phenomenon can be explained by the shape of our cost functional, where the cost of flying through risky zones depends not only on the probability level but on the time spent in risky zones, and thus flying faster is recommended. This might seem to contradict recomendations to not fly at high speed through turbulent air, but we address this issue with the v_{max} setting in Equation (20). The best solution, which does not cross zones where the probability field is above 2.5%, is the only one that does not feature these temporary airspeed increases.

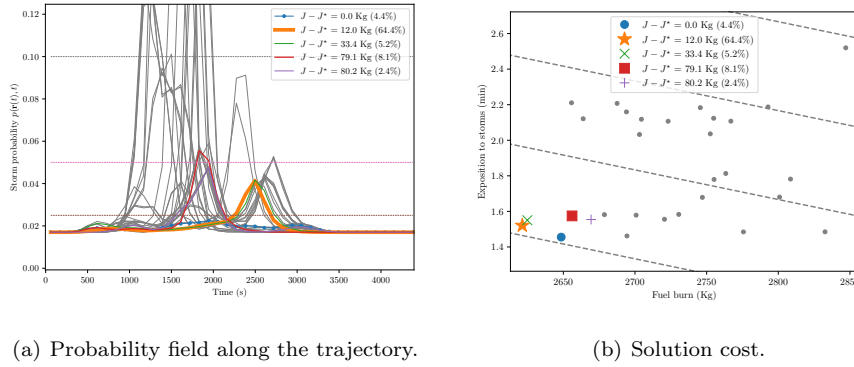(a) Probability field along the trajectory.

(b) Solution cost.

Figure 15: Probability of convection and cost. Fig a) The horizontal dashed lines represent the 2.5%, 5% and 10% probability levels. Fig b) The dashed lines represent equal total cost curves.

Figure 15(b) shows the two components of the objective value that is achieved by each trajectory. As discussed earlier, approximately 85% of the starting points lead to one of the five highlighted trajectories, which have close-to-optimal cost. A few of the remaining trajectories also have relatively small cost increases with respect to the optimum (50 - 200 Kg), manifesting mainly in the form of increased fuel burn. The second cluster has similar fuel consumption patterns, but the potential exposition to convective weather rises by about 50% (from 1.4 equivalent minutes to 2.1 equivalent minutes).

37

The information provided in Figure 15 could be presented to the decision-maker in order to assess the safety of the proposed trajectory, as Figure 15(a) represents the risk of encountering thunderstorms at each instant in the trajectory while Figure 15(b) illustrates the accumulated risk. Note that, in our approach, the event to avoid (because it would likely lead to an unplanned route deviation) is an encounter with a convective cell, not a instance of loss of separation with the convective cell, as we deal with uncertainty in an explicit form instead of employing pre-determined safety margins. Nevertheless, it is certainly possible to add additional safety margins to our approach by enlarging the deterministically nowcasted storms by the desired amount before the application of the model.


Figure 16: Distribution of the computational cost.

The distribution of the computational cost associated with the solution of the optimization problem (not including preprocessing) is shown in Figure 16. It can be observed that the NLP solver takes less than 30 seconds of CPU time to find the corresponding solution around 90% of the time. The computational cost of the algorithm can be thus considered as moderate. While the 10-40 seconds of

optimization time is useful for analysis and close to real-time performance, the multiple runs needed to obtain the alternative trajectories with the initialization scheme represent a limit in terms of practical implementation. Nevertheless, because runs are independent processes, it is possible to run them in parallel.

5. Conclusions

We have presented a data-driven methodology for estimating a probabilistic thunderstorm nowcasting based a existing, deterministic thunderstorm nowcasting product. We then combined it with a robust optimal control methodology for the design of avoidance trajectories under uncertainty on the motion of thunderstorms. It considers nonlinear aircraft dynamics, therefore generating realistic trajectories within a moderate computational cost. The initialization scheme allows the methodology to capture multiple local optima and identify the best one, as well as providing alternative trajectories. In a future concept of operations, the pilot or air traffic controller might prefer to employ these alternative options because of considerations not included in the optimization model, such as conflicts with other aircraft. In this way, multi-aircraft situations could be managed in a more computationally affordable way than a full n-aircraft optimal control problem, an approach that requires the consideration of $\mathcal{O}(n^2)$ potential interactions.

Possible extensions exist in the two main areas of contribution of the paper: First, the probabilistic model for the motion of the thunderstorms can be enhanced by the addition of other meteorological variables to the set of predictors. Additionally, the growth and decay processes could be considered to enhance the model. Finally, a statistical analysis of the relationship between meteorological conditions and flight plan deviations could be performed in order to create a more applicable indicator. Second, the optimal control formulation and the application of direct methods allow for direct and efficient extension of the methodology to consider additional constraints or dynamics in an efficient fashion; in particular, the methodology could be extended in future work to a full 4-D model, with movement in the vertical plane and not just the horizontal plane.

Acknowledgments

References

[1] A. J. Cook, G. Tanner, European airline delay cost reference values.

[2] D. González Arribas, M. F. Soler Arnedo, M. Sanjurjo Rivo, J. García-Heras Carretero, D. Sacher, U. Gelhardt, J. Lang, T. Hauf, J. Simarro, Robust Optimal Trajectory Planning under Uncertain Winds and Convective Risk, in: 5th ENRI Workshop on ATM/CNS, 2017.

[3] J. W. Wilson, N. A. Crook, C. K. Mueller, J. Sun, M. Dixon, Nowcasting Thunderstorms: A Status Report, Bull. Amer. Meteor. Soc. 79 (10) (1998) 2079–2099. doi:10.1175/1520-0477(1998)079<2079:ntasr>2.0.co;2.
URL https://doi.org/10.1175/1520-0477(1998)079<2079:ntasr>2.0.co;2

[4] J. E. Evans, E. R. Ducot, Corridor integrated weather system, Lincoln Laboratory Journal 16 (1) (2006) 59.

[5] J. Sun, M. Xue, J. W. Wilson, I. Zawadzki, S. P. Ballard, J. Onvlee-Hooimeyer, P. Joe, D. M. Barker, P.-W. Li, B. Golding, M. Xu, J. Pinto, Use of NWP for Nowcasting Convective Precipitation: Recent Progress and Challenges, Bull. Amer. Meteor. Soc. 95 (3) (2014) 409–426. arXiv:https://doi.org/10.1175/BAMS-D-11-00263.1, doi:10.1175/bams-d-11-00263.1.
URL https://doi.org/10.1175/bams-d-11-00263.1

---

[12] https://optmet.wordpress.com/

[6] Y. Hwang, A. J. Clark, V. Lakshmanan, S. E. Koch, Improved Nowcasts by Blending Extrapolation and Model Forecasts, Wea. Forecasting 30 (5) (2015) 1201–1217. doi:10.1175/waf-d-15-0057.1.
URL https://doi.org/10.1175/waf-d-15-0057.1

[7] M. M. Wolfson, W. J. Dupree, R. M. Rasmussen, M. Steiner, S. G. Benjamin, S. S. Weygandt, Consolidated storm prediction for aviation (CoSPA), in: 2008 Integrated Communications, Navigation and Surveillance Conference, IEEE, 2008. doi:10.1109/icnsurv.2008.4559190.
URL https://doi.org/10.1109/icnsurv.2008.4559190

[8] R. Osinski, F. Bouttier, Short-range probabilistic forecasting of convective risks for aviation based on a lagged-average-forecast ensemble approach, Met. Apps 25 (1) (2017) 105–118. doi:10.1002/met.1674.
URL https://doi.org/10.1002/met.1674

[9] M. Sauer, T. Hauf, C. Forster, Uncertainty Analysis of Thunderstorm Nowcasts for Utilization in Aircraft Routing, in: Fourth SESAR Innovation Days, Madrid, Spain, 2014.

[10] L. Han, J. Sun, W. Zhang, Y. Xiu, H. Feng, Y. Lin, A machine learning nowcasting method based on real-time reanalysis data, J. Geophys. Res. Atmos. 122 (7) (2017) 4038–4051. arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1002/2016JD025783, doi:10.1002/2016jd025783.
URL https://doi.org/10.1002/2016jd025783

[11] S. Xingjian, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, W.-c. Woo, Convolutional LSTM network: A machine learning approach for precipitation nowcasting, in: Advances in neural information processing systems, 2015, pp. 802–810.

[12] R. DeLaura, M. Robinson, M. Pawlak, J. Evans, Modeling convective weather avoidance in enroute airspace, in: 13th Conference on Aviation, Range and Aerospace Meteorology, 2008.

[13] R. DeLaura, B. Crowe, R. Ferris, J. Love, W. Chan, Comparing convective weather avoidance models and aircraft-based data, in: 89th Annual Meeting of the American Meteorological Society: Aviation, Range and Aerospace Meteorology Special Symposium on Weather–Air Traffic Impacts, Phoenix, AZ., 2009.

[14] M. Matthews, R. DeLaura, Evaluation of enroute convective weather avoidance models based on planned and observed flights, in: 14th American Meteorological Society Conference on Aviation, Range, and Aerospace Meteorology, 2010.

[15] K. Sheth, T. Amis, S. Gutierrez-Nolasco, B. Sridhar, D. Mulfinger, Development of a Probabilistic Convective Weather Forecast Threshold Parameter for Flight-Routing Decisions, Wea. Forecasting 28 (5) (2013) 1175–1187. doi:10.1175/waf-d-12-00052.1.
URL https://doi.org/10.1175/waf-d-12-00052.1

[16] S.-C. Wu, C. L. Gooding, A. E. Shelley, C. G. Duong, W. W. Johnson, Pilot convective weather decision making in en route airspace, in: International Congress of the Aeronautical Sciences (ICAS 2012), 2012.

[17] H. Erzberger, T. A. Lauderdale, Y.-C. Chu, Automated conflict resolution, arrival management, and weather avoidance for air traffic management, Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering 226 (8) (2011) 930–949. doi:10.1177/0954410011417347.
URL https://doi.org/10.1177/0954410011417347

[18] H. Erzberger, T. Nikoleris, R. A. Paielli, Y.-C. Chu, Algorithms for control of arrival and departure traffic in terminal airspace, Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering 230 (9) (2016) 1762–1779. doi:10.1177/0954410016629499.
URL https://doi.org/10.1177/0954410016629499

[19] R. Windhorst, M. Refai, S. Karahan, Convective weather avoidance with uncertain weather forecasts, in: 2009 IEEE/AIAA 28th Digital Avionics Systems Conference, IEEE, 2009, pp. 3.D.4–1–3.D.4–10. `doi:10.1109/dasc.2009.5347511`.
URL `https://doi.org/10.1109/dasc.2009.5347511`

[20] C. Taylor, S. Liu, C. Wanke, T. Stewart, Generating Diverse Reroutes for Tactical Constraint Avoidance, Journal of Air Transportation 26 (2) (2018) 49–59. `doi:10.2514/1.d0089`.
URL `https://doi.org/10.2514/1.d0089`

[21] H. K. Ng, S. Grabbe, A. Mukherjee, Design and Evaluation of a Dynamic Programming Flight Routing Algorithm Using the Convective Weather Avoidance Model, in: AIAA Guidance, Navigation, and Control Conference, American Institute of Aeronautics and Astronautics, 2009. `doi:10.2514/6.2009-5862`.
URL `https://doi.org/10.2514/6.2009-5862`

[22] B. Zhang, L. Tang, M. Roemer, Probabilistic Planning and Risk Evaluation Based on Ensemble Weather Forecasting, IEEE Trans. Automat. Sci. Eng. 15 (2) (2018) 556–566. `doi:10.1109/tase.2017.2648743`.
URL `https://doi.org/10.1109/tase.2017.2648743`

[23] M. Kamgarpour, V. Dadok, C. Tomlin, Trajectory generation for aircraft subject to dynamic weather uncertainty, in: 49th IEEE Conference on Decision and Control (CDC), IEEE, IEEE, 2010, pp. 2063–2068. `doi:10.1109/cdc.2010.5717889`.
URL `https://doi.org/10.1109/cdc.2010.5717889`

[24] W. Liu, I. Hwang, Probabilistic Aircraft Midair Conflict Resolution Using Stochastic Optimal Control, IEEE Trans. Intell. Transport. Syst. 15 (1) (2014) 37–46. `doi:10.1109/tits.2013.2274999`.
URL `https://doi.org/10.1109/tits.2013.2274999`

[25] S. Summers, M. Kamgarpour, J. Lygeros, C. Tomlin, A stochastic reach-avoid problem with random obstacles, in: Proceedings of the 14th international conference on Hybrid systems: computation and control - HSCC '11, ACM, ACM Press, 2011, pp. 251–260. doi:10.1145/1967701.1967738.
URL https://doi.org/10.1145/1967701.1967738

[26] D. González-Arribas, D. Hentzen, M. Sanjurjo-Rivo, M. Soler, M. Kamgarpour, Optimal Aircraft Trajectory Planning in the Presence of Stochastic Convective Weather Cells, in: 17th AIAA Aviation Technology, Integration, and Operations Conference, American Institute of Aeronautics and Astronautics, 2017. doi:10.2514/6.2017-3431.
URL https://doi.org/10.2514/6.2017-3431

[27] D. Hentzen, M. Kamgarpour, M. Soler, D. González-Arribas, On maximizing safety in stochastic aircraft trajectory planning with uncertain thunderstorm development, Aerosp. Sci. Technol. 79 (2018) 543–553. doi:10.1016/j.ast.2018.06.006.
URL https://doi.org/10.1016/j.ast.2018.06.006

[28] K. A. Emanuel, Atmospheric convection, Oxford University Press, 1994.

[29] G. J. Snyder, Introduction, Vol. 88, NYU Press, 2018.
URL https://doi.org/10.18574/nyu/9780814769867.003.0002

[30] Y.-L. Lin, Mesoscale Dynamics, Cambridge University Press, 2007. doi:10.1017/cbo9780511619649.
URL https://doi.org/10.1017/cbo9780511619649

[31] F. Molteni, T. Stockdale, M. Balmaseda, G. Balsamo, R. Buizza, L. Ferranti, L. Magnusson, K. Mogensen, T. Palmer, F. Vitart, The new ECMWF seasonal forecast system (System 4), Tech. rep. (2011).

[32] Y.-Y. Park, R. Buizza, M. Leutbecher, TIGGE: Preliminary results on comparing and combining ensembles, Q.J.R. Meteorol. Soc. 134 (637) (2008)

2029–2050. doi:10.1002/qj.334.

URL https://doi.org/10.1002/qj.334

[33] T. Iversen, A. Deckmyn, C. Santos, K. Sattler, J. B. Bremnes, H. Feddersen, I.-L. Frogner, Evaluation of 'GLAMEPS'—a proposed multimodel EPS for short range forecasting, Tellus A: Dynamic Meteorology and Oceanography 63 (3) (2011) 513–530. doi:10.1111/j.1600-0870.2010.00507.x.

URL https://doi.org/10.1111/j.1600-0870.2010.00507.x

[34] Y. Wang, M. Bellus, C. Wittmann, M. Steinheimer, F. Weidle, A. Kann, S. Ivatek-Šahdan, W. Tian, X. Ma, S. Tascu, E. Bazile, The Central European limited-area ensemble forecasting system: Aladin-laef, Q.J.R. Meteorol. Soc. 137 (655) (2011) 483–502. doi:10.1002/qj.751.

URL https://doi.org/10.1002/qj.751

[35] B. Øksendal, Stochastic Differential Equations, Springer Berlin Heidelberg, 2003. doi:10.1007/978-3-642-14394-6.

URL https://doi.org/10.1007/978-3-642-14394-6

[36] A. Gardi, R. Sabatini, S. Ramasamy, Multi-objective optimisation of aircraft flight trajectories in the ATM and avionics context, Progress in Aerospace Sciences 83 (2016) 1–36. doi:10.1016/j.paerosci.2015.11.006.

URL https://doi.org/10.1016%2Fj.paerosci.2015.11.006

[37] E. Trélat, Optimal Control and Applications to Aerospace: Some Results and Challenges, J Optim Theory Appl 154 (3) (2012) 713–758. doi:10.1007/s10957-012-0050-5.

URL https://doi.org/10.1007/s10957-012-0050-5

[38] A. V. Rao, A survey of numerical methods for optimal control, Advances in the Astronautical Sciences 135 (1) (2009) 497–528. doi:10.1002/oca.2114.

[39] J. T. Betts, Practical Methods for Optimal Control and Estimation Using Nonlinear Programming, Vol. 19, Society for Industrial and Applied Mathematics, 2010. doi:10.1137/1.9780898718577.
URL https://doi.org/10.1137/1.9780898718577

[40] P. E. Gill, W. Murray, M. A. Saunders, SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization, SIAM Rev. 47 (1) (2005) 99–131. doi:10.1137/s0036144504446096.
URL https://doi.org/10.1137/s0036144504446096

[41] A. Wächter, L. T. Biegler, On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming, Math. Program. 106 (1) (2005) 25–57. doi:10.1007/s10107-004-0559-y.
URL https://doi.org/10.1007/s10107-004-0559-y

[42] E. Gallo, F. Navarro, A. Nuic, M. Iagaru, Advanced Aircraft Performance Modeling for ATM: Bada 4.0 Results, in: 2006 ieee/aiaa 25TH Digital Avionics Systems Conference, IEEE, 2006, pp. 1–12. doi:10.1109/dasc.2006.313660.
URL https://doi.org/10.1109/dasc.2006.313660

[43] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, M. Diehl, CasADi: A software framework for nonlinear optimization and optimal control, Math. Prog. Comp.doi:10.1007/s12532-018-0139-4.
URL https://doi.org/10.1007/s12532-018-0139-4

[44] HSL. A collection of Fortran codes for large scale scientific computation. http://www.hsl.rl.ac.uk/.

[45] I. M. Ross, M. Karpenko, A review of pseudospectral optimal control: From theory to flight, Annual Reviews in Control 36 (2) (2012) 182–197. doi:10.1016/j.arcontrol.2012.09.002.
URL https://doi.org/10.1016/j.arcontrol.2012.09.002

[46] A. Franco, D. Rivas, A. Valenzuela, Minimum-Fuel Cruise at Constant Altitude with Fixed Arrival Time, Journal of Guidance, Control, and Dynamics 33 (1) (2010) 280–285. doi:10.2514/1.46465.
URL https://doi.org/10.2514/1.46465