



Article submitted to journal

**Subject Areas:**

Computational Mechanics, Computer Modelling and Simulation

**Keywords:**

Trajectory Prediction, Probabilistic Machine Learning, Functional Data Analysis, Gaussian Process Emulators, Monotonicity

**Author for correspondence:**

Nick Pepper

e-mail: [npepper@turing.ac.uk](mailto:npepper@turing.ac.uk)

# A Probabilistic Model for Aircraft in Climb using Monotonic Functional Gaussian Process Emulators

Nick Pepper<sup>1</sup>, Marc Thomas<sup>2</sup>, George De Ath<sup>3</sup>, Enrico Olivier<sup>3</sup>, Richard Cannon<sup>2</sup>, Richard Everson<sup>3</sup> and Tim Dodwell<sup>4</sup>

<sup>1</sup> The Alan Turing Institute, The British Library, London, UK.

<sup>2</sup> NATS

<sup>3</sup> Department of Computer Science, University of Exeter, Exeter, UK.

<sup>4</sup> digiLab, Exeter, UK.

Ensuring vertical separation is a key means of maintaining safe separation between aircraft in congested airspace. Aircraft trajectories are modelled in the presence of significant epistemic uncertainty, leading to discrepancies between observed trajectories and the predictions of deterministic models, hampering the task of planning to ensure safe separation. In this paper a probabilistic model is presented, for the purpose of emulating the trajectories of aircraft in climb and bounding the uncertainty of the predicted trajectory. A monotonic, functional representation exploits the spatio-temporal correlations in the radar observations. Through the use of Gaussian Process Emulators, features that parameterise the climb are mapped directly to functional outputs, providing a fast approximation, while ensuring that the resulting trajectory is monotonic. The model was applied as a probabilistic digital twin for aircraft in climb and baselined against the Base of Aircraft Data (BADA), a deterministic model widely used in industry. When applied to an unseen test dataset, the probabilistic model was found to provide a mean prediction that was 20.56% more accurate, as measured by the mean absolute error, with data-driven credible intervals that were 9.54% sharper.

## 1. Introduction

Trajectory prediction (TP) plays an important role in the safe management of aircraft and is used by air traffic controllers to inform decision-making by predicting aircraft trajectories and detecting potential conflicts [1]. To ensure safety, air traffic controllers strive to maintain either longitudinal or vertical separation of aircraft [2]. In this context vertical TP, the prediction of an aircraft's altitude with time, is especially important as controllers must devise plans that maintain separation at all times. State-of-the-art TP methods project the future path of an aircraft through models of the flight mechanics of aircraft, in which an estimate of the present state of the aircraft is used as an initial condition for the numerical solution of a set of equations approximating the physics governing the aircraft's flight (see, e.g. [3–5]). However, there is a high level of epistemic uncertainty inherent in this approach due to: the simplifications necessarily made by the model; an uncertain knowledge of the aircraft's state; lack of knowledge of the pilot's intentions [6]; and the unknown influence of environmental effects on the aircraft trajectory [7]. As a consequence a mismatch between the predictions of physics-based TP methods and the actual path followed by aircraft can be observed, especially during climbs and descents [8,9].

The observed mismatch between physics-based methods and real trajectories has motivated the application of machine learning-based methods to TP. In some respects TP is a very amenable problem for machine learning due to the large quantity of radar observations available to train models [10–13]. In recent years, methods based on Neural Networks [14,15] have been proposed for TP. However, there are several challenges facing machine learning methods for TP that must be addressed. Firstly, given that high levels of epistemic uncertainty contribute so significantly to model-mismatch in state-of-the-art equations-based models, a probabilistic approach would appear to be essential. Next generation TP models must be able to efficiently handle uncertainties and to clearly express the uncertainty in the model predictions [7,16]. This is the motivation for approaches based on Sequential Monte Carlo sampling [17] and Gaussian Mixture Models in the TP literature [1,18]. Indeed, several recent reviews of Machine Learning methods have expressed this same point, that moving away from deterministic black-box machine learning models is necessary for the massive industrial application of machine learning [19,20], particularly in risk-averse industries such as aeronautics. Uncertainty estimates in TP for climbing aircraft are not generally data-driven. For instance, the left panel of Figure 1 illustrates an envelope created by two runs of a deterministic TP model with minimum and maximum nominal mass, where the nominal mass refers to the aircraft mass expected by the model. However, this is an over-conservative approach. What is desired from a probabilistic TP model is illustrated on the right panel of Figure 1: a data-driven credible interval that trades off some conservatism in favour of more realistically representing the uncertainty.

In parallel with the continued refinement of TP methods, there has been a drive in recent years to explore the application of Artificial Intelligence (AI) to emulate the role of an Air Traffic Controller (ATCO) that has been motivated by the maturity of dynamic optimisation and reinforcement learning [21,22]. In this context, the probabilistic approach to TP appears to be particularly useful. During training, an AI agent's plan is evaluated within a Digital Twin of an airspace. These plans must be robust to variabilities in the trajectories of aircraft, requiring a probabilistic digital twin. This offers another area of application for a probabilistic TP model [23].

The second challenge facing ML methods is to guarantee that predicted trajectories satisfy physical constraints, particularly when testing on unseen data. In the case of TP, there is a requirement that these trajectories are achievable for the performance envelope of an aircraft. Simultaneously, there is a qualitative constraint that predicted trajectories correspond to modes of aircraft operation that are observed in the real-world. In the case of aircraft in climb, a reasonable expectation is that the altitude of an aircraft will increase monotonically with time, until a target altitude is reached. Enforcing a monotonicity constraint is straightforward if the ML approach corrects the parameters of an equation-based model. For instance, Alligier et al. [24,25] proposed a ML model to better predict the mass and speed intent of climbing aircraft, unknown parameters




Figure 1: An uncertainty bound for an aircraft climbing between flight levels 200 and 300, obtained from runs of a deterministic TP code with maximal and minimal mass, compared to the observed radar data (left). An illustration of the desired data-driven bounds, that better reflects the observed trajectories (right).

which are then fed into the Base of Aircraft Data (BADA) physics-based model. However, there is a computational cost associated with the numerical solution of BADA, especially in a probabilistic approach that might require multiple model evaluations. A fast approximation to BADA is required, where model inputs are mapped directly to trajectories. However, guaranteeing that such fast approximations will themselves be monotonic is challenging. Some progress has been made in this regard in the setting of pointwise data, for example recent work with Gaussian Process Emulators (GPEs) has found correlation functions that correspond to a monotonic mean function [26,27]. However, little attention has been paid in the ML literature to fast approximations with monotonic functional outputs.

Finally, a point prediction approach to TP, for instance by using a neural network to predict future states of an aircraft, may be limited as it does not fully exploit the spatio-temporal correlations available in the radar observations [28]. Instead, a functional approach may be a more natural way to express the predicted trajectory. This is the motivating philosophy in the papers of Nicol et al. [29,30], in which functional Principal Component Analysis (fPCA) is used to analyse aircraft trajectories. In this work we propose a method for vertical TP using a functional Gaussian Process approach in which we marry ideas from Functional Data Analysis (FDA) with Gaussian Process Emulators. The model produces a functional estimate of a trajectory that is informed by the physics of the problem, that is, the trajectory is constrained to be monotonic. By exploiting the posterior variance, the predictive uncertainty of the method can be expressed by sampling the posterior distribution. The novel contributions of this paper may be summarised as follows:

- A generative model for aircraft in climb that exploits tempo-spatial correlations in the trajectory data
- Posterior samples from the model are guaranteed to be monotonic, a necessary condition for them to be considered physically plausible
- Gaussian Process Emulators are used to generate monotonic samples conveniently using standard choices of kernel functions

In the following section we describe the probabilistic model, before demonstrating its application to a dataset of real flight data in Section 3, where the model is baselined against BADA, a deterministic TP code used widely in industry.

## 2. A Probabilistic model for climbing aircraft

We propose a probabilistic model that can express the flight level,  $f$ , as a function of time,  $t$ , for an aircraft that is cleared to climb by an Air Traffic Control Officer (ATCO). Flight levels refer to the altitude at standard air pressure, expressed in hundreds of feet. A probabilistic model for the climb is denoted  $f(t|\mathbf{x})$ , where  $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^{n_x}$  represents a set of  $n_x$  features that parameterise the climb. In what follows these features are considered to be time independent, consisting of the instructions issued to the aircraft by the ATCO and the available data pertaining to the aircraft's state when the command was issued. The model is trained using a set of radar observations containing  $n_d$  monotonically increasing trajectories,  $D = \{(\mathbf{x}^{(i)}, f^{(i)}(t))\}$ ,  $i = 1 \dots n_d$ . One radar sweep takes approximately 6 seconds, while the duration of a typical climb in  $D$  is of the order of 5-10 minutes. The method consists of two parts: firstly, trajectories in  $D$  are described with a functional, monotonic representation that is parameterised by a set of hyper-parameters,  $\mathbf{y} \in \mathcal{Y} \subseteq \mathbb{R}^{n_y}$ ; and secondly, a set of  $n_y$  Gaussian Process Emulators are defined to accomplish the probabilistic mapping  $\mathbf{x} \rightarrow \mathbf{y}$ .

### (a) Monotonic representation of functional data

Functional Principal Components Analysis (fPCA) is a popular tool in FDA for performing dimensionality reduction on functional data. In fPCA a function is expressed as a weighted sum of orthonormal basis functions,  $\phi_i(t)$ , and a mean function,  $\mu(t)$ :

$$f(t) \approx \mu(t) + \sum_{i=1}^{n_c} \alpha_i \phi_i(t), \quad (2.1)$$

where  $\alpha \in \mathbb{R}^{n_c}$  are referred to as the Principal Component scores (with  $\alpha_i$  denoting the  $i$ -th score) and  $n_c$  is the number of Principal Components [31]. fPCA has proven to provide successful representations of functional data, however, the process of finding orthogonal basis vectors is purely data-driven and is oblivious to the physical constraints on the data generating process. For instance, in the targeted application of vertical TP it is expected that  $f(t)$  will be a smooth monotonic function, with constraints imposed on its derivative with respect to time,  $\partial_t f$ , by the performance of the aircraft. Since they are orthogonal, the fPCA modes,  $\phi_i(t)$ , are oscillatory. These basis functions are chosen such that they commit the smallest  $L_2$  error for each  $n_c$  among all possible bases.

In theory an infinite number of modes are required to represent monotonic functions. However, in practice the summation is truncated and it is therefore possible for the set of Principal Component scores,  $\alpha$ , to correspond to a trajectory that is not monotonic even if all the trajectories in  $D$  are themselves monotonic [32]. For this reason we, instead, employ the monotonic representation of functional data from Ramsay [33]. Provided that  $f$  satisfies the conditions that:

- $\log \partial_t f$  is differentiable;
- $\partial_t \log \partial_t f = \partial_t^2 f / \partial_t f$  is Lebesgue square integrable;

then such a monotonic function may be represented using an integral form:

$$f(t) = \beta_0 + \beta_1 \int_{\tau_0}^t \exp \int_{\tau_0}^s w(u) du ds, \quad (2.2)$$

where  $\beta_0$  and  $\beta_1$  are coefficients to be determined and  $w(t)$  represents a square integrable function such that  $w = \partial_t^2 f / \partial_t f$ .  $\tau_0$  represents the time when the manoeuvre begins. The second condition requires that the ratio of the curvature of the trajectory to its slope are bounded. Practically, satisfying these conditions requires the trajectory to be strictly monotonic (i.e.  $\partial_t f > 0$ ) and that  $\partial_t^2 f \rightarrow 0$  only if  $\partial_t f \rightarrow 0$ . Given that we expect  $\partial_t f$  to be continuous for an aircraft in climb, these conditions are considered reasonable for aircraft trajectories. The conditions make it difficult to




Figure 2: Plot of a trajectory (red) and its reconstructions (blue) using the monotonic representation proposed here with varying numbers of Fourier modes (left). Associated weighting functions  $w(t)$  (right).

model an aircraft leveling out at its target flight level, because aircraft often level out between radar measurements and, therefore, have observed first derivative that is discontinuous. In what follows the trajectory is truncated at the target flight level, instead of modelling the abrupt change in climb rate. This approach is reasonable given that the target flight level is known *a priori*, and the period of the levelling out (shorter than 6s) is much shorter than the period of the entire climb. Inspired by the work of Shin et al. [34], in which fPCA is used to represent this function, we cast  $w(t)$  as a Fourier series:

$$w(t) = a_0 + \sum_{i=1}^{n_w} \mathbf{a}_i \cos(2\pi i t) + \mathbf{b}_i \sin(2\pi i t), \quad (2.3)$$

where the set of coefficients  $\mathbf{a}, \mathbf{b} \in \mathbb{R}^{n_w}$ ,  $\beta_1$ , and  $a_0$  are determined through Stochastic Gradient Descent (SGD) for each trajectory in  $D$ , in which the residual sum of squares (RSS) loss between the observations and the functional representation is minimised.  $\mathbf{a}_i$  and  $\mathbf{b}_i$  denote the  $i^{\text{th}}$  components of  $\mathbf{a}$  and  $\mathbf{b}$ , respectively. In what follows we enforce the initial condition  $f(\tau_0 = 0) = f_i$ , where  $f_i$  represents the flight level of the first observation after the clearance to climb was issued. As a consequence  $\beta_0 = f_i$ . Figure 2 illustrates the reconstruction of a normalised trajectory in  $D$  using this formulation and the associated  $w(t)$ . Note that there is some rigidity in the representation, in the example shown increasing the number of Fourier modes from 3 to 20 reduces the error in the representation of the data by 62%, which we denote  $\epsilon_{\mathcal{L}}$ . However, we do not expect  $\epsilon_{\mathcal{L}} \rightarrow 0$  as  $n_w \rightarrow \infty$  because of the rigidity.

The rationale for the use of the Fourier series is to select a basis that is guaranteed to be orthogonal. This addresses a disadvantage of the method proposed by Shin, where a transformation for the B-splines used as basis functions must be found to orthogonalise them (which we note has the effect of adding oscillations into the resulting basis). A further benefit is that  $w(t)$  does not need to be integrated numerically as analytical expressions are available. The Ramsay framework provides a bijective representation of the functional data in  $D$ . The developments in Appendix A detail the SGD procedure used to find the set of optimal parameters in the representation,  $\hat{\mathbf{y}} = [\hat{\beta}_1, \hat{a}_0, \hat{\mathbf{a}}, \hat{\mathbf{b}}]^\top$ , that, when repeated for each of the trajectories in  $D$ , yields the set  $D_y = \{\hat{\mathbf{y}}^{(i)}\}$ ,  $i = 1, \dots, n_d$  (with  $n_y = 2n_w + 2$ ).

### (i) Dimensional Reduction with Principal Component Analysis

Having described the  $n_d$  climbs in  $D$  with the monotonic framework of Ramsay, a model is desired to learn the mapping between the features  $\mathbf{x}$  and the corresponding  $2n_w + 2$  parameters in this framework. To simplify this mapping, a projection of the data is performed using a

Principal Component Analysis (PCA) on  $D_y$ . Having performed this projection, the coefficients may be expressed using the finite sum:

$$\mathbf{y} = \begin{bmatrix} \beta_1 \\ a_0 \\ \mathbf{a} \\ \mathbf{b} \end{bmatrix} \approx \sum_{k=1}^{n_c} \boldsymbol{\alpha}_k \boldsymbol{\psi}_k \quad (2.4)$$

where  $\boldsymbol{\alpha} \in \mathbb{R}^{n_c}$  is a vector of Principal Component (PC) scores and  $\boldsymbol{\psi}_k \in \mathbb{R}^{2n_w+2} \subseteq \mathcal{Y}$  is the  $k^{\text{th}}$  Principal Component.  $n_c \in [1, 2n_w + 2]$  is determined through an adaptive process, described in Appendix B. This projection simplifies the mapping from  $\mathbf{x}$  because the data is now expressed in a lower-dimensional, statistically decorrelated basis (assuming  $n_c < 2n_w + 2$ ). There are now two levels of approximation in the model: one in the finite sum in (2.3) and another by the finite sum in (2.4). Expressing the PC scores as a probabilistic function of  $\mathbf{x}$ , we can rewrite equations (2.2) and (2.3) to include a dependence on the features parameterising each climb:

$$f(t|\mathbf{x}) = f_i + \sum_{k=1}^{n_c} \boldsymbol{\alpha}_k(\mathbf{x}) \boldsymbol{\psi}_{1,k} \int_{\tau_0}^t \exp \int_{\tau_0}^s w(u|\mathbf{x}) du ds, \quad (2.5)$$

with

$$w(t|\mathbf{x}) = \sum_{j=2}^{2n_w+2} \sum_{k=1}^{n_c} \boldsymbol{\alpha}_k(\mathbf{x}) \boldsymbol{\psi}_{j,k} \boldsymbol{\theta}_j(t), \quad (2.6)$$

where  $\boldsymbol{\psi}_{j,k}$  refers to the  $j^{\text{th}}$  element of the  $k^{\text{th}}$  Principal Component and the vector  $\boldsymbol{\theta}(t) \in \mathbb{R}^{2n_w+1}$  contains the Fourier modes, evaluated at time  $t$ , i.e.:

$$\boldsymbol{\theta}(t) = [1, \cos(2\pi t), \dots, \cos(2\pi n_w t), \sin(2\pi t), \dots, \sin(2\pi n_w t)]^\top. \quad (2.7)$$

Projecting each trajectory in  $D$  onto this new basis generates the set  $D_\alpha = \{\boldsymbol{\alpha}^{(i)}\}$ ,  $i = 1, \dots, n_d$ . The sub-set  $D_x = \{\mathbf{x}^{(i)}\}$ ,  $i = 1, \dots, n_d$  and  $D_\alpha$  are used as the training data for a set of independent Gaussian Process Emulators (GPEs), that are outlined in the next section.

## (b) Gaussian Process Emulators

Having found a monotonic functional representation of each trajectory in  $D$ , an interpolation function,  $\Psi$ , is developed for each component of  $\boldsymbol{\alpha}$ :

$$\hat{\boldsymbol{\alpha}}_k(\mathbf{x}) = \Psi_k(\mathbf{x}) + \epsilon_k, \quad k = 1, \dots, n_c, \quad (2.8)$$

where  $\epsilon \sim \mathcal{N}(\mathbf{0}, \text{diag}(\sigma_\epsilon))$  is assumed to be independent, identically distributed Gaussian noise with variance  $\sigma_\epsilon^2$ . A set of GPEs are used to find the interpolated PCA scores  $\hat{\boldsymbol{\alpha}}$  for a test parameter point  $\mathbf{x}^* \subseteq \mathcal{X}$ . We provide a brief summary of GPEs here, the interested reader is referred to the canonical reference books of Rasmussen and Williams [35] and Santner et al. [36]. A Gaussian Process prior is assumed over the regression functions, i.e.

$$\begin{bmatrix} A_k \\ \hat{\boldsymbol{\alpha}}_k(\mathbf{x}^*) \end{bmatrix} \sim \mathcal{N} \left( \begin{bmatrix} \mu(X) \\ \mu(\mathbf{x}^*) \end{bmatrix}, \begin{bmatrix} K(X, X) & K(\mathbf{x}^*, X)^T \\ K(\mathbf{x}^*, X) & K(\mathbf{x}^*, \mathbf{x}^*) \end{bmatrix} \right), \quad (2.9)$$

where  $\mu$  is a function representing the mean of the process (in what follows we set  $\mu = 0$ ). The matrix,  $K$ , reflects the covariances of the training data, which is organised into a  $n_x \times n_d$  matrix of inputs  $X = [\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(n_d)}]^\top$  with outputs  $A_k = [\boldsymbol{\alpha}_k^{(1)}, \boldsymbol{\alpha}_k^{(2)}, \dots, \boldsymbol{\alpha}_k^{(n_d)}]$ . Elements of the

covariance matrix are given by:

$$K_{ij}(X, X) = k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) + \sigma_{\epsilon_k}^2 \delta_{ij}, \quad (2.10)$$

where  $k(\cdot)$  represents the covariance function and  $\delta$  the Kronecker delta. In this work squared exponential kernels were used for the covariance function. The hyperparameters for the GPE are inferred from the data, at a cost  $O(n_d^3)$ . To mitigate this cost when  $n_d$  is large (e.g.  $n_d > 10^3$ ), we use a stochastic variational regression for the GP hyperparameters. We refer the interested reader to Blei et al. [37] and Hensman et al. [38] for more details on this. A consequence of assuming a Gaussian Process prior is that the posterior predictive density is also Gaussian:

$$\hat{\boldsymbol{\alpha}}_k | X, A_k, \mathbf{x}^* \sim \mathcal{N}(\hat{\mu}_k, \hat{\sigma}_k), \quad (2.11)$$

where

$$\begin{aligned} \hat{\mu}_k &= K(\mathbf{x}^*, X)^\top K(X, X)^{-1} A_k, \\ \hat{\sigma}_k^2 &= K(\mathbf{x}^*, \mathbf{x}^*) - K(\mathbf{x}^*, X)^\top K(X, X)^{-1} K(X, \mathbf{x}^*). \end{aligned} \quad (2.12)$$

This is advantageous for the targeted application because the posterior distribution can be resampled for negligible cost. For a given set of features,  $\mathbf{x}^*$ ,  $n_{mc}$  Monte Carlo samples can be drawn, giving the set  $\{\boldsymbol{\alpha}^{(i)}, \dots, \boldsymbol{\alpha}^{(n_{mc})}\}$ . When passed through the monotonic framework of (2.5) and (2.6) this yields the family of trajectories  $\{f^{(i)}(t|\mathbf{x}^*)\}$ ,  $i = 1, \dots, n_{mc}$ . The main steps of the model training may be summarised as follows:

**Algorithm 1 (Training of a monotonic, functional Gaussian Process Emulator model for aircraft in climb):**

Inputs: Training dataset,  $D$

Outputs: Gaussian Process Emulators  $N(\hat{\mu}_k, \hat{\sigma}_k)$ ,  $k = 1, \dots, n_c$

- (i) Obtain parameters of the Ramsay monotonic representation ( $\mathbf{y} = [\beta_1, a_0, \mathbf{a}, \mathbf{b}]^\top$ ) for each trajectory in  $D$ , using SGD. Collect them as  $D_y$  (A 1)
- (ii) Perform PCA on  $D_y$  (2.4)
- (iii) Truncate PC representation (Appendix B)
- (iv) Project onto the basis of retained PCs, yielding set of PC scores  $D_\alpha$  (2.5) & (2.6)
- (v) Extract sub-sequence of clearance instructions  $D_x$  from  $D$
- (vi) Train set of GPEs to perform the probabilistic mapping  $\mathbf{x} \rightarrow \boldsymbol{\alpha}$  using stochastic variational regression (2.8)

Having described the main features of the algorithm, in the following section we demonstrate its application to a dataset of radar observations from within a sector of UK airspace.

### 3. Application to real trajectory data

In this section we demonstrate the application of the probabilistic model to a dataset containing 10,827 aircraft trajectories. This data was harvested from ATC data pertaining to a sector of UK airspace between January and February 2018. This sector contains one of the UK's busiest international airports, and as a consequence, most aircraft in the dataset are climbing to reach their cruising altitude. From the available data the most suitable features were identified by industry experts that were expected to parameterise the climbs:

- (i) The requested change in flight level,  $\Delta f$
- (ii) The initial flight level,  $f_i$
- (iii) The aircraft indicated airspeed,  $v_{ias}$






Figure 3: Visualisation of the ground tracks of the climbs in the dataset, where each line indicates the route of a single flight in the dataset.

The three features are collected in the input vector  $\mathbf{x} = [\Delta f, f_i, v_{ias}]^T$ . Ideally, industry experts would recommend working with the true airspeed, rather than the indicated airspeed. Such a conversion is possible provided that weather data and a model for the measurement uncertainty in  $v_{ias}$  is available in the dataset. In what follows we use  $v_{ias}$  but note that, should it be available, it would be possible to incorporate this additional data without altering the fundamentals of the monotonic functional method demonstrated here. Examples of parameters that are understood to significantly affect the trajectory but are not included in the dataset, because they are either unknown or proprietary, include:

- The effect of variations in weather on the aircraft
- Mass of the aircraft
- True airspeed of the aircraft
- Aircraft performance settings

Rather than attempt to model the effects of these parameters, or assume probability distributions for them, the effect of this epistemic uncertainty is absorbed within the posterior distribution of the GPEs. In this section we test whether such a data-driven approach can provide a credible fast approximation for the vertical trajectory of an aircraft. The data-driven uncertainty envelope is compared against one inferred from multiple runs of BADA sampled from a distribution for the aircraft mass. BADA, the deterministic TP model used as a baseline, is an energy-based model that relates to the geometrical, kinematic and kinetic aspects of the aircraft motion, allowing the aircraft performances and trajectory to be predicted. BADA is a deterministic model that is widely used by industry. Given that the output of BADA is sensitive to the mass variation of the aircraft, we contrast the probabilistic method proposed with a simple probabilistic model using BADA, in which the aircraft mass is sampled from a naive probability distribution. This simple probabilistic model is discussed further in subsection (b). Figure 3 illustrates the ground tracks of the climbs in the dataset.



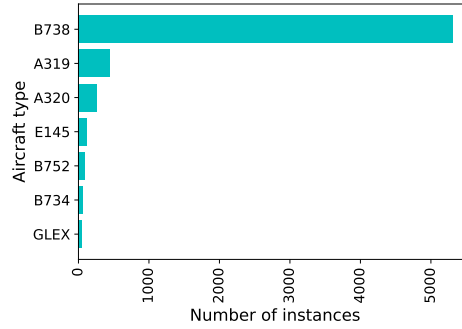


Figure 4: Trajectories in the dataset by aircraft type (types with more than 50 instances plotted)

### (a) Comparison with BADA

The outputs of the probabilistic model were first baselined against a single, deterministic run of BADA, using the nominal mass of the aircraft. For consistency, trajectories from a single aircraft type, the B738, and a single operator were used. As can be seen from Figure 4, this aircraft type dominated the dataset: 77.7% of the aircraft in the dataset are B738s and of these 92.5% are operated by one carrier. From these trajectories, a dataset of 4612 climbs, where the aircraft received a clearance to climb by 80 or more flight levels, were extracted for use in analysis. The distribution of  $\Delta f$  is plotted in the top left panel of Figure 5. The top right panel indicates the discrepancies in the arrival time at the target flight level between BADA and the dataset, normalised by the predicted length of the climb. The resulting distribution is Gaussian-like with a positive mean, indicating that in general BADA over-predicts the arrival time. The bottom panel of Figure 5 plots this discrepancy as a function of the target flight level, with the shading of the markers indicating  $\Delta f$  for each climb. From around flight level 300 onwards the markers generally lie in the positive half-plane. This would suggest that the sector specific speed profiles do not exactly match the BADA calibrated speed profiles. While the parameters in the BADA model are calibrated globally, some local deviation between observed and predicted trajectories can sometimes be observed, as has been examined in studies such as Alligier and Gianazza [25] and Sun et al. [43].

The probabilistic model was tested on this dataset of climbs through a five-fold cross-validated procedure. The dataset consisted of data from 47 days of operations. The probabilistic model was trained five times, with a different test set of consecutive days' of data held out each time. In this way, each of the 4612 climbs was included in a test set. In what follows, the results of the five test sets are reported together. A set of  $n_{mc} = 100$  posterior samples were drawn from the probabilistic model for each test trajectory.

A set of performance indicators were used to assess the relative skilfulness of the TP models. One of the selected indicators was the predicted time for each trajectory to achieve the target flight level, which we denote  $t_a$ . Taken together, a set of posterior samples provide a probabilistic forecast for  $t_a$  that could be scored using the Continuous Ranked Probability Score (CRPS) [44,45]:

$$CRPS^{(i)} = \int_{-\infty}^{\infty} \left( F^{(i)}(t) - \mathbb{1}(t - t_a^{(i)}) \right)^2 dt, \quad (3.1)$$

where  $F^{(i)}(t)$  is the CDF for  $t_a$  from the  $n_{mc}$  Monte Carlo samples;  $\mathbb{1}(\cdot)$  the Heaviside step function; and  $t_a^{(i)}$  the observed arrival time [39]. For a deterministic forecast (i.e. a single run of BADA), the CRPS score is equivalent to the  $L_2$  error. By taking the ratio of the CRPS score for the probabilistic model and the equivalent score for BADA, the relative skilfulness of the model




Figure 5: Top left: distribution of the requested change in flight levels in the dataset. Top right: Distribution of errors in the arrival time (BADA-data), normalised by the predicted length of the climb, according to BADA. Bottom: Normalised arrival time error as a function of the target flight level. Markers are coloured by  $\Delta f$  for each climb.

may be expressed as a dimensionless skill score:

$$S^{(i)} = 1 - \frac{CRPS^{(i)}}{(t_a^{(i)} - t_{a,BADA}^{(i)})^2}, \quad (3.2)$$

where  $S^{(i)}$  represents the skill score for the  $i^{\text{th}}$  trajectory in the test set and  $t_{a,BADA}^{(i)}$  the arrival time for that trajectory as predicted by BADA (with nominal mass value). Positive values of  $S$  indicate that the probabilistic model is more skilful. The left panel of Figure 6 is a histogram displaying the skill scores, with the results from the 5 test datasets amalgamated. The median skill score was found to be +0.3474, indicating that the probabilistic model was more skilful than the deterministic run of BADA for this performance indicator.

The offset in the predictions at the arrival time,  $\Delta z$ , was also used as a performance indicator. From an ATC perspective this is a useful quantity because it is connected to the flight level occupancy of the aircraft. For the  $i^{\text{th}}$  trajectory in the test dataset, with observed arrival time,  $t_a$  and estimated arrival time  $\hat{t}_a$ , this quantity is denoted:

$$\Delta z^{(i)} = \begin{cases} (f_i + \Delta f) - f(t_a | \mathbf{x}^{(i)}) & \text{if } t_a \leq \hat{t}_a, \\ (f_i + \Delta f) - f^{(i)}(\hat{t}_a) & \text{if } t_a > \hat{t}_a \end{cases}. \quad (3.3)$$

Lower values of  $\Delta z$  indicate that a model has more accurately predicted the flight level occupancy at the arrival time. The right panel of Figure 6 displays two (normalised) histograms for  $\Delta z$ , with  $\Delta z$  calculated for each trajectory in the test dataset for BADA (green) and for the mean of the probabilistic model (blue). As can be seen, these discrepancies are comparable, and the mean discrepancy across all 4612 test flights is tabulated in Table 1. This Table also tabulates the Mean Absolute Error (MAE) between the means of the models and the observed trajectory. The




Figure 6: Left: Histogram displaying the distribution of skill scores for the test dataset. The vertical line indicates the median skill. Right: Histogram displaying normalised histograms of  $\Delta z$  for the probabilistic model (blue) versus BADA (green).

mean prediction of the probabilistic model offered a 20.56% improvement in accuracy for this test dataset.

Figure 7 visualises the output of the probabilistic model for a random selection of eight trajectories from the first test dataset. The red line indicates the observed radar data, the green line the solution of BADA for the requested climb, and the blue line the mean prediction of the generative model. The variance in the arrival time at intermediate flight levels is used to construct a  $2\sigma$  credible region for the probabilistic model. Note that in some instances, such as the second plotted test flight, the data-driven credible interval is sufficiently tight that it contains the observed trajectory but not the deterministic run of BADA. The skill scores and sharpness of each test flight are tabulated in Table 2. In the next section, the probabilistic model is compared against a probabilistic implementation of BADA and the two metrics used to assess the credible intervals are introduced.

### (b) Comparison with BADA using assumed aircraft mass distribution

Having compared the probabilistic model against BADA, run using the nominal mass value for the aircraft, in this section we make a comparison between the proposed model and a simple probabilistic implementation of BADA. Given the sensitivity of the BADA predictions to the mass of the aircraft, we formulate a simple probabilistic model for BADA in which the mass of the aircraft is sampled from a uniform distribution  $U(m_{min}, m_{max})$ , where  $m_{min}$  and  $m_{max}$  are the minimum and maximum masses associated with that aircraft in the BADA model. As before, we evaluate the two probabilistic models using the arrival time and offset in flight levels at the arrival time as performance indicators. In this case, the relative skilfulness of the proposed probabilistic model, baselined by the probabilistic BADA prediction, is denoted:

$$S_{pb}^{(i)} = 1 - \frac{CRPS^{(i)}}{CRPS_{pb}^{(i)}}, \quad (3.4)$$

where  $S_{pb}$  refers to the relative skilfulness, baselined by the probabilistic implementation of BADA, with associated CRPS score  $CRPS_{pb}$ . The left panel of Figure 9 indicates the relative skilfulness for each trajectory in the test set. The median skill score of +0.8917 was higher than when the probabilistic model was baselined against the deterministic model evaluation. This is likely due to the much higher variance in the sampled trajectories from the probabilistic BADA, as can be seen in Figure 8. The CRPS score penalises overly-conservative forecasts. In Figure 8 the  $2\sigma$  credible region from the probabilistic model is compared to the credible region from the probabilistic BADA model, with the data-driven bound from the probabilistic model being

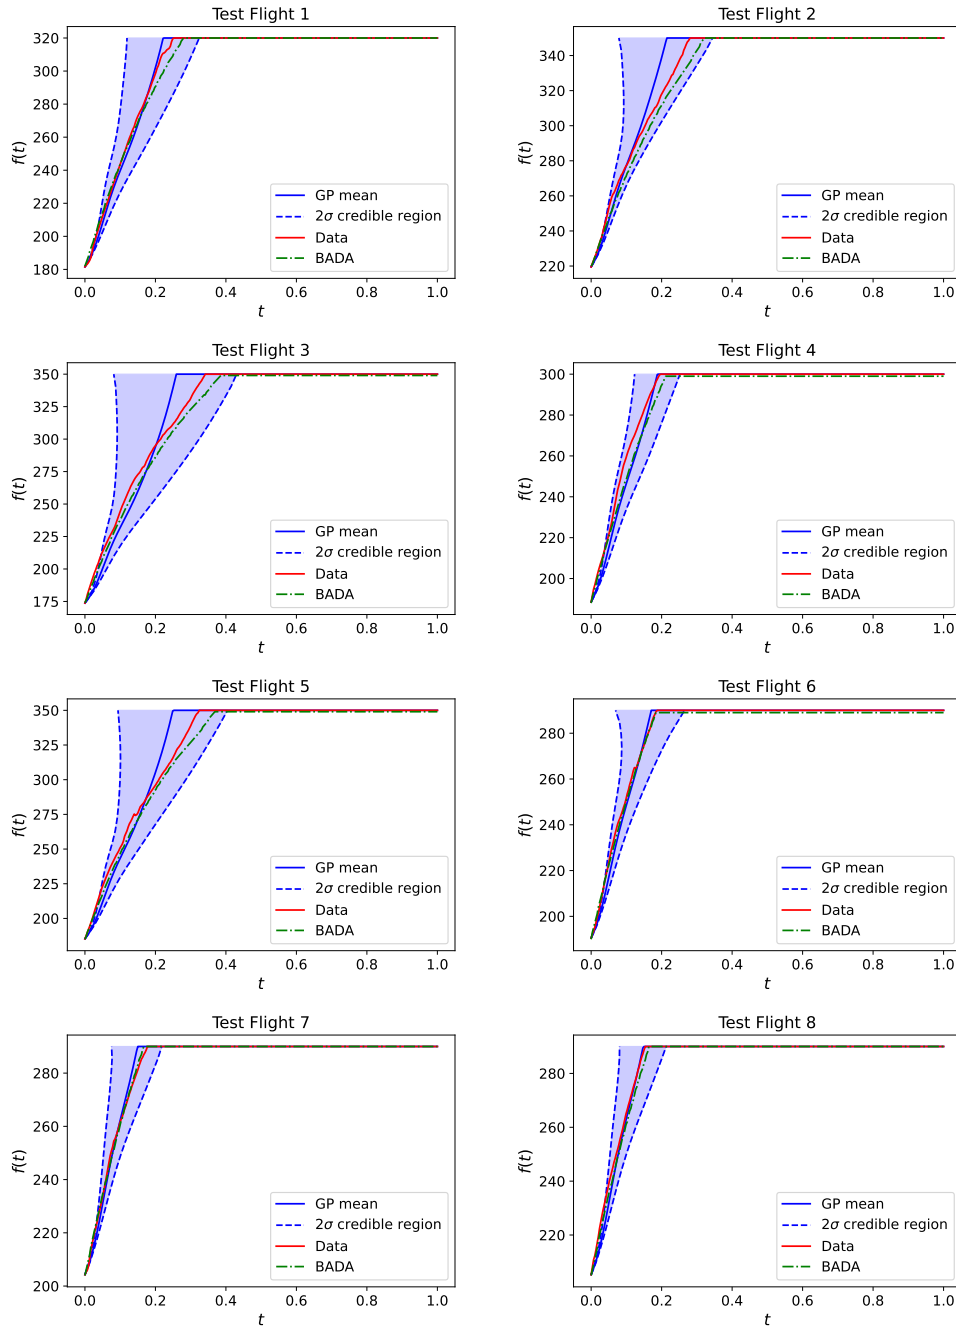


Figure 7: Radar data for an unseen trajectory (red), compared to the predicted trajectory by BADA (green), and the mean of the probabilistic digital twin (blue) for eight trajectories in the test dataset. The dashed lines indicate the  $2\sigma$  credible region from 100 posterior samples. Note that the credible regions are drawn with respect to arrival time, rather than flight level.

noticeably tighter. These confidence bounds were assessed using the Root Mean Squared Error of Calibration (RMSEC) [40–42] and the sharpness. The RMSEC quantifies how well calibrated the two probabilistic TP models are, while the sharpness quantifies the concentration of the uncertainty. Ideally, we wish for a data-driven model to provide a sharper forecast than BADA, subject to calibration. Here we find that the data-driven confidence intervals are 9.54% sharper, with the total calibration error decreasing by 11.49%. To determine these values, a set of 20 intermediate flight levels were chosen for each test flight and the mean and variance of the arrival time were computed.

Interestingly, the mean of the posterior samples was found to be more accurate than the nominal BADA mass, as quantified by the MAE. This is because the reference mass used in the run of BADA with nominal parameters is greater than the mean of the uniform prior for the aircraft mass. Samples from the prior will in general climb faster than the nominal run of BADA. Recalling from Figure 5 that BADA with nominal parameters tended to over-estimate the length of climb, it is to be expected that the mean of these samples will be more accurate.

Finally, the right panel of Figure 9 displays the normalised histograms for  $\Delta z$ . Note that in this case  $\Delta z$  was calculated for each of the posterior samples, rather than the mean estimate, hence the chart for the probabilistic model differs from that in Figure 6. The data-driven posterior samples were found to have a much lower average discrepancy.

Table 1: The performance of the probabilistic model, bench-marked against BADA for various metrics, including: the Mean Absolute Error (MAE), median relative skill score for predicting the end of climb (for the probabilistic model), flight level discrepancy at the end of climb ( $\Delta z$ ), and Root Mean Squared Error of Calibration (RMSEC). All quantities are averaged over the five folds. Arrows indicate whether higher or lower values of that metric are preferable.

Method	MAE ↓	$S$ ↑	$\Delta z$ ↓	RMSEC ↓	Sharpness ↓
Prob. Model	<b>1.9600</b>	-	14.9851 (mean) 19.0264 (post. samp.)	<b>0.2173</b>	<b>0.0237</b>
BADA (nom.)	2.4674	0.3474	<b>13.5139</b>	-	-
BADA (prob.)	2.0631	0.8917	18.6779	0.2455	0.0262

Table 2: Skill scores for the eight plotted test flights and the relative sharpness of the methods

Test flight	$S$ ↑		Relative sharpness ↓
	BADA (nom.)	BADA (prob.)	
1	<b>0.3621</b>	0.9212	0.8741
2	-0.0518	0.8404	0.7320
3	-0.0376	0.8464	1.0845
4	0.3298	0.9469	<b>0.6695</b>
5	0.0357	0.8480	0.8434
6	-1.9266	0.9400	0.9471
7	-0.6283	0.8943	0.8336
8	0.3552	<b>0.9505</b>	0.7810

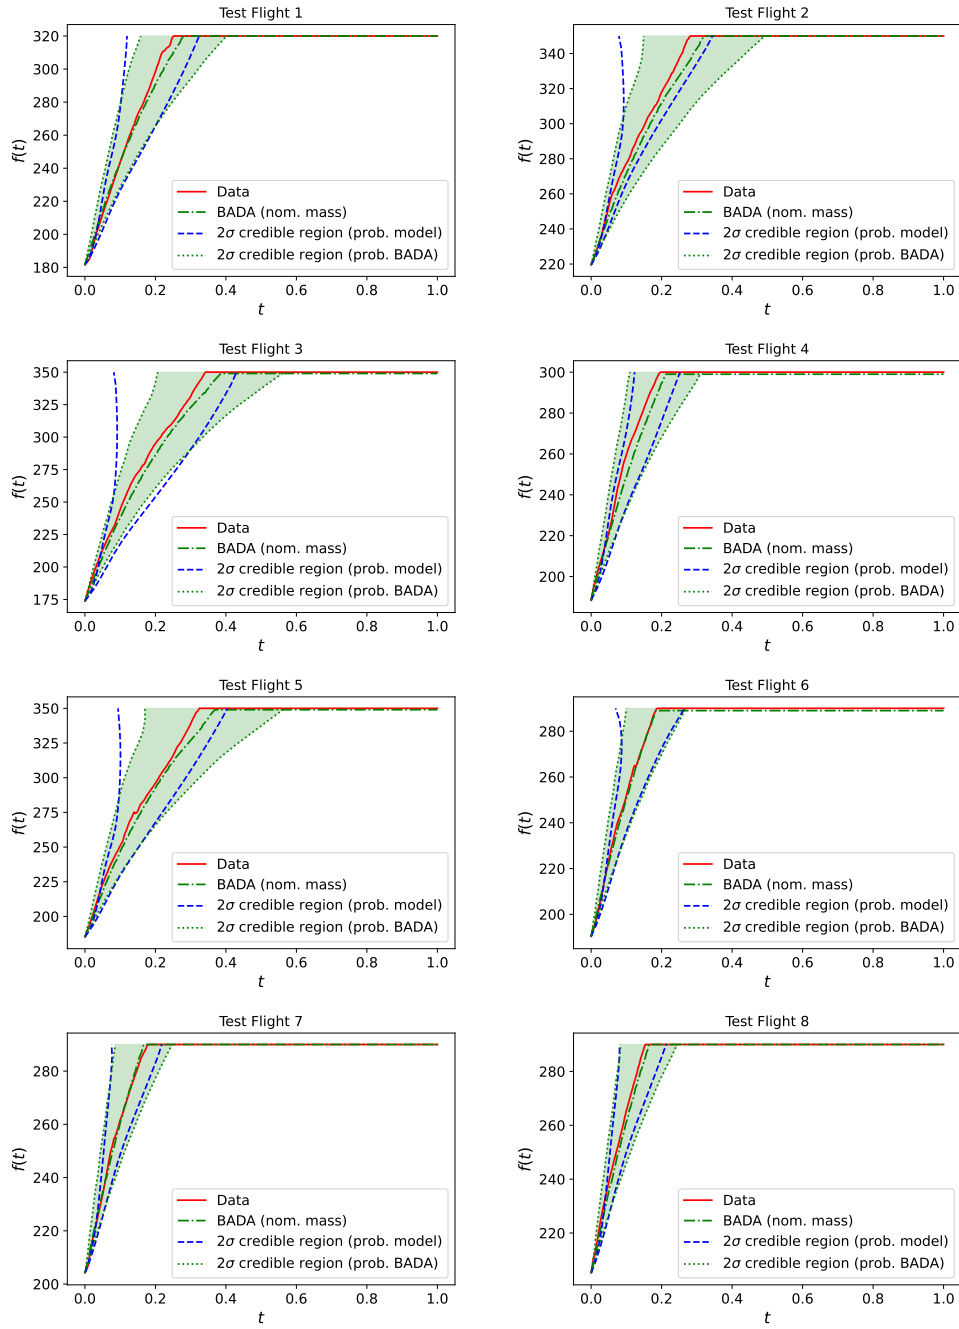


Figure 8: Visual of the predicted trajectories for the same eight test points using the probabilistic implementation of BADA. The observed trajectory (red), is compared against the BADA run using nominal mass (green). The credible interval of the model is also plotted (dashed green) versus that of the proposed model (dashed blue).




Figure 9: Left: Histogram displaying the distribution of skill scores for the probabilistic model, baselined by a probabilistic implementation of BADA. Right: Normalised histograms for the discrepancies in the flight level at the end of climb for the samples.

## 4. Conclusion

Accurate trajectory prediction for an aircraft in climb is difficult due to the presence of significant epistemic uncertainty. We have proposed a probabilistic model that employs Gaussian Process Emulators to give a fast approximation to a trajectory that is functional and monotonic. A functional approach allows tempo-spatial correlations in the training data to be exploited. In this paper the proposed method has been baselined against a deterministic TP model, widely used in industry, in addition to a probabilistic version that exploits the available knowledge of the aircraft's mass. Using test datasets containing unseen data, it was found that the proposed model outperformed both the deterministic and probabilistic versions of BADA, with performance quantified by several metrics including the relative skill of the model in predicting the arrival time of an aircraft at its target flight level.

In this paper we have restricted the application of the model to TP for climbing aircraft, although we note that the method is general and could be applied to any class of problem where a fast approximation for monotonic functional outputs is required. Finally, in its presented form the method is completely data-driven. However, we note that a deterministic model could be used to inform the model by acting as a mean function for the GPEs. In such an instance, it would be necessary to find a mapping such that solutions of the deterministic model could be expressed in the coefficient space of the Ramsay framework.

**Authors' Contributions.** N.P. Investigation and Writing. M.T. Data curation and Methodology. G.D. and E.O. Methodology and Software. R.C., R.E., and T.D. Conceptualization and Supervision. All authors reviewed and edited the manuscript.

**Competing Interests.** M.T. and R.C. are employed by NATS, the air traffic service provider for the United Kingdom.

**Funding.** The work described in this article is primarily funded by the grant "EP/V056522/1 : Advancing Probabilistic Machine Learning to Deliver Safer, More Efficient and Predictable Air Traffic Control" (aka Project Bluebird), and EPSRC Prosperity Partnership between NATS, Turing, Exeter and Cambridge. T.D. acknowledges funds from his UKRI Turing AI Fellowship 2TAFFP\100007.

## References

1. Barratt, S., Kochenderfer, M. & Boyd, S. Learning Probabilistic Trajectory Models of Aircraft in Terminal Airspace From Position Data. *IEEE Transactions On Intelligent Transportation Systems*. 20 pp. 3536–3545 (2019)
2. Paielli, R., Erzberger, H., Chiu, D. & Heere, K. Tactical Conflict Alerting Aid for Air Traffic Controllers. *Journal Of Guidance, Control, and Dynamics*. 32 pp. 184–193 (2009)



3. Chatterji, G., Sridhar, B. & Bilimoria, K. En-route flight trajectory prediction for conflict avoidance and traffic management. *Guidance, Navigation, And Control Conference.*, <https://arc.aiaa.org/doi/abs/10.2514/6.1996-3766>
4. Chatterji, G. Short-term trajectory prediction methods. *Guidance, Navigation, And Control Conference And Exhibit.*, <https://arc.aiaa.org/doi/abs/10.2514/6.1999-4233>
5. Nuic, A., Poles, D. & Mouillet, V. BADA: An advanced aircraft performance model for present and future ATM systems. *International Journal Of Adaptive Control And Signal Processing.* **24**, 850–866 (2010), <https://onlinelibrary.wiley.com/doi/abs/10.1002/acs.1176>
6. Bastas, A., Kravaris, T. & Vouras, G. Data driven aircraft trajectory prediction with deep imitation learning. *ArXiv Preprint ArXiv:2005.07960*. (2020)
7. Pang, Y., Zhao, X., Yan, H. & Liu, Y. Data-driven trajectory prediction with weather uncertainties: A Bayesian deep learning approach. *Transportation Research Part C: Emerging Technologies.* **130** pp. 103326 (2021)
8. Baklacioglu, T. & Cavcar, M. Aero-propulsive modelling for climb and descent trajectory prediction of transport aircraft using genetic algorithms. *The Aeronautical Journal.* **118**, 65–79 (2014)
9. Thipphavong, D., Schultz, C., Lee, A. & Chan, S. Adaptive algorithm to improve trajectory prediction accuracy of climbing aircraft. *Journal Of Guidance, Control, And Dynamics.* **36**, 15–24 (2013)
10. De Leege, A., Paassen, M. & Mulder, M. A machine learning approach to trajectory prediction. *AIAA Guidance, Navigation, And Control (GNC) Conference.* pp. 4782 (2013)
11. Wang, Z., Liang, M. & Delahaye, D. Short-term 4d trajectory prediction using machine learning methods. *Proc. SID.* pp. 1–10 (2017)
12. Wu, X., Yang, H., Chen, H., Hu, Q. & Hu, H. Long-term 4D trajectory prediction using generative adversarial networks. *Transportation Research Part C: Emerging Technologies.* **136** pp. 103554 (2022)
13. Hernández, A., Casado Magaña, E. & Berna, A. Data-driven Aircraft Trajectory Predictions using Ensemble Meta-Estimators. *2018 IEEE/AIAA 37th Digital Avionics Systems Conference (DASC).* pp. 1–10 (2018)
14. Zhao, Z., Zeng, W., Quan, Z., Chen, M. & Yang, Z. Aircraft trajectory prediction using deep long short-term memory networks. *CICTP 2019.* pp. 124–135 (2019)
15. Shi, Z., Xu, M., Pan, Q., Yan, B. & Zhang, H. LSTM-based flight trajectory prediction. *2018 International Joint Conference On Neural Networks (IJCNN).* pp. 1–8 (2018)
16. Liu, W. & Hwang, I. Probabilistic Trajectory Prediction and Conflict Detection for Air Traffic Control. *Journal Of Guidance, Control, And Dynamics.* **34**, 1779–1789 (2011), <https://doi.org/10.2514/1.53645>
17. Lymperopoulos, I. & Lygeros, J. Sequential Monte Carlo methods for multi-aircraft trajectory prediction in air traffic management. *International Journal Of Adaptive Control And Signal Processing.* **24**, 830–849 (2010)
18. Paek, H., Lee, K. & Vela, A. En-route Arrival Time Prediction using Gaussian Mixture Model. (2020)
19. Hüllermeier, E. & Waegeman, W. Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods. *Machine Learning.* **110**, 457–506 (2021,3), <http://dx.doi.org/10.1007/s10994-021-05946-3>
20. Psaros, A., Meng, X., Zou, Z., Guo, L. & Karniadakis, G. Uncertainty Quantification in Scientific Machine Learning: Methods, Metrics, and Comparisons. (2022)
21. Brittain, M. & Wei, P. Autonomous air traffic controller: A deep multi-agent reinforcement learning approach. *ArXiv Preprint ArXiv:1905.01303*. (2019)
22. Ghosh, S., Laguna, S., Lim, S., Wynter, L. & Poonawala, H. A Deep Ensemble Method for Multi-Agent Reinforcement Learning: A Case Study on Air Traffic Control. *Proceedings Of The International Conference On Automated Planning And Scheduling.* **31**, 468–476 (2021,5), <https://ojs.aaai.org/index.php/ICAPS/article/view/15993>
23. Dodwell, T., Awad, E., Cannon, R., Carvell, B., De Ath, G., Everson, R., Gabasova, E., Girolami, M., Kamalaruban, P., Mei, P., Pace, A., Pepper, N., Thomas, M., Weller, A., Wordingham, F. Developing intelligent machines for air traffic control. *Nature: Machine Intelligence.*
24. Alligier, R., Gianazza, D. & Durand, N. Learning the aircraft mass and thrust to improve the ground-based trajectory prediction of climbing flights. *Transportation Research Part C: Emerging Technologies.* **36** pp. 45–60 (2013), <https://www.sciencedirect.com/science/article/pii/S0968090X13001708>

25. Alligier, R. & Gianazza, D. Learning aircraft operational factors to improve aircraft climb prediction: A large scale multi-airport study. *Transportation Research Part C: Emerging Technologies*. **96** pp. 72–95 (2018)
26. Riihimäki, J. & Vehtari, A. Gaussian processes with monotonicity information. *Proceedings Of The Thirteenth International Conference On Artificial Intelligence And Statistics*. pp. 645–652 (2010)
27. Ustyuzhaninov, I., Kazlauskaitė, I., Ek, C. & Campbell, N. Monotonic gaussian process flows. *International Conference On Artificial Intelligence And Statistics*. pp. 3057–3067 (2020)
28. Benavides, J., Kaneshige, J., Sharma, S., Panda, R. & Steglinski, M. Implementation of a Trajectory Prediction Function for Trajectory Based Operations. *AIAA Atmospheric Flight Mechanics Conference*, <https://arc.aiaa.org/doi/abs/10.2514/6.2014-2198>
29. Jarry, G., Delahaye, D., Nicol, F. & Feron, E. Aircraft atypical approach detection using functional principal component analysis. *Journal Of Air Transport Management*. **84** pp. 101787 (2020)
30. Nicol, F. Functional principal component analysis of aircraft trajectories. *ISIATM 2013, 2nd International Conference On Interdisciplinary Science For Innovative Air Traffic Management*. (2013)
31. Ramsay, J. & Silverman, B. *Functional Data Analysis*. (Springer New York, 2006), <https://books.google.co.uk/books?id=REzuyz>
32. Han, K., Hadjipantelis, P., Wang, J., Kramer, M., Yang, S., Martin, R. & Müller, H. Functional principal component analysis for identifying multivariate patterns and archetypes of growth, and their association with long-term cognitive development. *PloS One*. **13**, e0207073 (2018)
33. Ramsay, J. Estimating smooth monotone functions. *Journal Of The Royal Statistical Society: Series B (Statistical Methodology)*. **60**, pp. 365–375 (1998), <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/1467-9868.00130>
34. Shin, Y., Zhou, L. & Ding, Y. Joint estimation of monotone curves via functional principal component analysis. *Computational Statistics & Data Analysis*. **166** (2022), <https://ideas.repec.org/a/eee/csdata/v166y2022ics0167947321001778.html>
35. Rasmussen, C. & Williams, C. *Gaussian Processes for Machine Learning*. (The MIT Press, 2006)
36. Santner, T., Williams, B. & Notz, W. *The design and analysis of computer experiments*. (Springer, 2003)
37. Blei, D., Kucukelbir, A. & McAuliffe, J. Variational Inference: A Review for Statisticians. *Journal Of The American Statistical Association*. **112**, pp. 859–877 (2017,4), <https://doi.org/10.1080>
38. Hensman, J., Matthews, A. & Ghahramani, Z. Scalable Variational Gaussian Process Classification. (arXiv, 2014), <https://arxiv.org/abs/1411.2005>
39. Gneiting, T. & Raftery, A. Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal Of The American Statistical Association*. **102**, pp. 359–378 (2007), [10.1198/016214506000001437](https://doi.org/10.1198/016214506000001437)
40. Chung, Y., Char, I., Guo, H., Schneider, J. & Neiswanger, W. Uncertainty Toolbox: an Open-Source Library for Assessing, Visualizing, and Improving Uncertainty Quantification. (ArXiv, 2021), <https://arxiv.org/abs/2109.10254>
41. Tran, K., Neiswanger, W., Yoon, J., Zhang, Q., Xing, E. & Ulissi, Z. Methods for comparing uncertainty quantifications for material property predictions. *Machine Learning: Science And Technology*. **1**, 025006 (2020)
42. Kuleshov, V., Fenner, N. & Ermon, S. Accurate Uncertainties for Deep Learning Using Calibrated Regression. *CoRR*. **abs/1807.00263** (2018), <http://arxiv.org/abs/1807.00263>
43. Sun, J., Ellerbroek, J. & Hoekstra, J. Modeling aircraft performance parameters with open ADS-B data. *12th USA/Europe Air Traffic Management Research and Development Seminar*. (2017)
44. Matheson, J.E. & Winkler, R.L. Scoring rules for continuous probability distributions. *Management science*. **22** (1976)
45. Hersbach, H. Decomposition of the continuous ranked probability score for ensemble prediction systems. *Weather and Forecasting*. **15** (2000)
46. Duchi, J., Hazan, E. & Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization.. *Journal Of Machine Learning Research*. **12** (2011)
47. Zhang, J. Machine Learning With Feature Selection Using Principal Component Analysis for Malware Detection: A Case Study. (arXiv, 2019), <https://arxiv.org/abs/1902.03639>

## A. Parameter estimation using SGD

Trajectory data in the training set is collected from radar observations of an aircraft and is therefore discretised according to the number of radar ‘blips’ during the climb or descent.  $D$

is more realistically expressed as  $D = \{\mathbf{x}^{(i)}, f^{(i)}(t_j), j = 1 \dots n_b^{(i)}\}$  where  $i = 1 \dots n_d$  (there is no guarantee that each sampled trajectory will have the same number of radar observations). For the  $i^{\text{th}}$  trajectory we wish to estimate the parameters in the monotonic representation in (2.2) and (2.3). This can be cast as the optimisation problem:

$$\min_{\beta_1, a_0, \mathbf{a}, \mathbf{b}} \mathcal{L}_i(\beta_1, a_0, \mathbf{a}, \mathbf{b}), \quad (\text{A } 1)$$

where we define the loss function for the  $i^{\text{th}}$  trajectory,  $\mathcal{L}_i$ , as the residual sum of squares (RSS) loss between the observations and the functional representation, i.e.:

$$\mathcal{L}_i(\beta_1, a_0, \mathbf{a}, \mathbf{b}) = \sum_{j=1}^{n_b^{(i)}} \{f_j^{(i)} - f(t_j | \beta_1, a_0, \mathbf{a}, \mathbf{b})\}^2. \quad (\text{A } 2)$$

Analytic expressions for the derivatives of the loss function are available:

$$\frac{\partial \mathcal{L}_i}{\partial f_j} = -2\{f_j^{(i)} - f_j\}, \quad (\text{A } 3)$$

where we have defined  $f_j = f(t_j | \beta_1, a_0, \mathbf{a}, \mathbf{b})$  for ease of notation. Using the chain rule we can find derivatives for each of the parameters:

$$\frac{\partial \mathcal{L}_i}{\partial \beta_1} = \sum_{j=1}^{n_b^{(i)}} \frac{\partial \mathcal{L}_i}{\partial f_j} \frac{\partial f_j}{\partial \beta_1} = -2 \sum_{j=1}^{n_b^{(i)}} \{f_j^{(i)} - f_j\} \int_{\tau_0}^{t_j} \exp \int_{\tau_0}^s w(u) du ds \quad (\text{A } 4)$$

$$\frac{\partial \mathcal{L}_i}{\partial a_0} = \sum_{j=1}^{n_b^{(i)}} \frac{\partial \mathcal{L}_i}{\partial f_j} \frac{\partial f_j}{\partial a_0} = -2\beta_1 \sum_{j=1}^{n_b^{(i)}} \{f_j^{(i)} - f_j\} \{\exp(t_j - \tau_0) - 1\} \quad (\text{A } 5)$$

$$\frac{\partial \mathcal{L}_i}{\partial \mathbf{a}_k} = \sum_{j=1}^{n_b^{(i)}} \frac{\partial \mathcal{L}_i}{\partial f_j} \frac{\partial f_j}{\partial \mathbf{a}_k} = -2\beta_1 \sum_{j=1}^{n_b^{(i)}} \{f_j^{(i)} - f_j\} \int_{\tau_0}^{t_j} \exp \int_{\tau_0}^s \cos(2\pi k u) du ds \quad (\text{A } 6)$$

$$\frac{\partial \mathcal{L}_i}{\partial \mathbf{b}_k} = \sum_{j=1}^{n_b^{(i)}} \frac{\partial \mathcal{L}_i}{\partial f_j} \frac{\partial f_j}{\partial \mathbf{b}_k} = -2\beta_1 \sum_{j=1}^{n_b^{(i)}} \{f_j^{(i)} - f_j\} \int_{\tau_0}^{t_j} \exp \int_{\tau_0}^s \sin(2\pi k u) du ds \quad (\text{A } 7)$$

Adagrad is a gradient based optimisation algorithm used to update the parameters with a variable learning rate [46]. At the  $m^{\text{th}}$  iteration the set of parameters,  $\mathbf{y} = \{\beta_1, a_0, \mathbf{a}, \mathbf{b}\}$  is adjusted according to:

$$\mathbf{y}_{l,m} = \mathbf{y}_{l,m-1} - \frac{\eta}{\sqrt{G_{mll}} + \rho} \frac{\partial \mathcal{L}_i}{\partial \mathbf{y}_{l,m-1}}, \quad (\text{A } 8)$$

where  $G_m$  is a diagonal matrix, in which the element  $G_{mll}$  contains the sums of the squares of the past gradients for the  $l^{\text{th}}$  parameter. The constant  $\rho$  is a smoothing constant used to ensure numerical stability when the gradients are small, while the learning rate  $\eta \sim \mathcal{O}(10^{-2})$  is set by the user. Here we used  $\eta = 0.02$ .

SGD is run to convergence for each trajectory in  $D$ , defined by a stopping criterion for the relative improvement of  $\mathcal{L}$ , yielding the set of parameters  $D_y = \{\hat{\mathbf{y}}^{(i)}\}$ ,  $i = 1, \dots, n_d$ . In what has been presented the search for the parameters is unconstrained, although we note that performance constraints of the aircraft, if known, could be expressed through a bound on  $w$ . Such constraints could be expressed through a feasibility function  $g(w)$ , yielding the modified optimisation problem:

$$\min_{\beta_1, a_0, \mathbf{a}, \mathbf{b}} \mathcal{L}_i(\beta_1, a_0, \mathbf{a}, \mathbf{b}) \quad (\text{A } 9)$$

$$\text{such that } g(w | \beta_1, a_0, \mathbf{a}, \mathbf{b}, \mathbf{x}_{per}) > 0,$$

where  $\mathbf{x}_{per}$  collects the additional performance related data required to evaluate  $g(\cdot)$ .

## B. Sub-algorithm to truncate Principal Component representation

Section 2(a) discusses how PCA is used to rotate  $D_y$  into a basis with statistically independent components. A further advantage of using the PCA representation is that the PCs are ranked in terms of the explained variance, allowing PCs that do not contribute significantly to be neglected. Typically, this is done through an analysis of the explained variance ratio (see, e.g. [47]). However, in this case it is difficult to quantify how variance in  $\mathbf{y}$  scales with variance in the generated trajectories when passed through the monotonic framework. For this reason, we propose a simple adaptive method for selecting  $n_c$ , based on the error associated with reconstructing the trajectories in  $D$  after they have been projected into the PC basis. We denote this quantity  $\epsilon_r$ , calculated through:

$$\epsilon_r = \sum_{i=1}^{n_d} \sum_{j=1}^{n_b^{(i)}} \left( f^{(i)}(t_j) - f(t_j | \mathbf{x}^{(i)}, \mathbf{y}) \right)^2, \quad (\text{A } 1)$$

where

$$\mathbf{y} = \sum_{k=1}^{n_c} \alpha_k \psi_k. \quad (\text{A } 2)$$

We remove PCs until a terminal condition on  $\epsilon_r$  is met. PC removal is terminated when the reconstruction error exceeds 10% of it's value at  $n_c = 2n_f + 2$ . This is illustrated in Figure 10 for the first training dataset. For the sake of comparison, we plot the error introduced in the parameter estimation,  $\epsilon_{\mathcal{L}}$ . This is defined as  $\epsilon_{\mathcal{L}} = \sum_{i=1}^{n_d} \mathcal{L}_i$  and in this case  $\epsilon_r \ll \epsilon_{\mathcal{L}}$ .




Figure 10: Plot displaying variation in  $\epsilon_r$  and  $\epsilon_{\mathcal{L}}$  (scaled by  $n_d$  and number of radar blips) with number of principal components. The vertical dashed line indicates the  $n_c$  chosen by the adaptive process.