

RESEARCH ARTICLE

From industry-wide parameters to aircraft-centric on-flight inference: Improving aeronautics performance prediction with machine learning

Florent Dewez¹, Benjamin Guedj^{1,2,*}  and Vincent Vandewalle^{1,3}

¹MODAL – MOdels for Data Analysis and Learning, Inria, Lille - Nord Europe Research Centre, France

²Department of Computer Science, Centre for Artificial Intelligence, University College London, London, United Kingdom

³CHU Lille, ULR 2694 Evaluations des technologies de santé et des pratiques médicales, University of Lille, Lille, France

*Corresponding author. E-mail: benjamin.guedj@inria.fr

Received: 11 May 2020; **Revised:** 28 July 2020; **Accepted:** 14 September 2020

Keywords: Aeronautics performance prediction; aircraft performance; aircraft performance monitoring; machine learning; statistical modeling

Abstract

Aircraft performance models play a key role in airline operations, especially in planning a fuel-efficient flight. In practice, manufacturers provide guidelines which are slightly modified throughout the aircraft life cycle via the tuning of a single factor, enabling better fuel predictions. However, this has limitations, in particular, they do not reflect the evolution of each feature impacting the aircraft performance. Our goal here is to overcome this limitation. The key contribution of the present article is to foster the use of machine learning to leverage the massive amounts of data continuously recorded during flights performed by an aircraft and provide models reflecting its actual and individual performance. We illustrate our approach by focusing on the estimation of the drag and lift coefficients from recorded flight data. As these coefficients are not directly recorded, we resort to aerodynamics approximations. As a safety check, we provide bounds to assess the accuracy of both the aerodynamics approximation and the statistical performance of our approach. We provide numerical results on a collection of machine learning algorithms. We report excellent accuracy on real-life data and exhibit empirical evidence to support our modeling, in coherence with aerodynamics principles.

Impact Statement

Current airline operations in both flight preparation (on-ground) and flight management (in-air) are mainly based on the performance of an aircraft. For instance, a trajectory is set before the flight and managed in-air by the Flight Management System using the manufacturer's performance model. This numerical model is calibrated during in-service period using monitoring systems developed by manufacturers. However, the calibration is based on the tuning of a single parameter, and this overly simplified modeling leads to a lack of precision in optimizing fuel consumption. In this paper, we propose performance models that take into account real flight conditions and switch from industry-wide to aircraft-centric calibration of relevant parameters. To do this, we use massive collections of in-air data recorded by the Quick Access Recorder, which reflect the actual behavior of the aircraft. We then resort to machine learning algorithms to learn sufficiently accurate models from this data to infer the actual performance of the aircraft. The present paper describes our overall approach and its application to

predicting the lift and drag coefficients. Bounds for the prediction errors are provided to assess the accuracy of the models. We aim at a twofold impact: (a) improve on the inference of in-flight parameters to optimize trajectories (e.g., regarding fuel consumption) and (b) provide a principled data-centric modeling approach which could be replicated in other intensive data-generating industries.

1. Introduction

Efficient aircraft operations require the knowledge of current performance of the aircraft. The performance depends on many aerodynamic and engines features which evolve throughout the life cycle of the aircraft. For instance, the degradation of the engines may modify the aircraft specific range up to 1.3% per year if no engine replacement is carried out (Airbus, 2002). Further aircraft physical features may change overtime due to accumulations of impurities on the surface (leading to an increase of the drag), rough or deformed surfaces, damaged seals, and so on (see for instance Airbus, 2001). In particular, performance may differ from an aircraft to another one of the same type due to different cycles, flight hours, or maintenance.

Given this deterioration, monitoring systems have been developed by the manufacturers (Krajčec et al., 2015). Those systems measure the difference between expected behavior obtained through heavy numerical simulations and wind tunnel tests, and data recorded during very specific in-flight conditions. Given this difference, most manufacturers propose to modify a performance index to adjust the theoretical fuel consumption to the current one. This leads to more accurate fuel predictions when planning the flight operations. However, according to Krajčec et al. (2015), such monitoring systems have limitations, for example, the incapacity of dissociating the influence of aerodynamic or engines features on the evolution of performance, since it is only described by the performance index.

In view of this, a flexible methodology designed to build models for features which reflect the current performance of an aircraft would be relevant to improve flight planning operations in the end. Such a methodology is proposed in the present paper.

The underlying idea of our approach is that the real performance of an aircraft should be reflected by its data recorded in recent flights. Here, we consider data from the Quick Access Recorder (QAR) which contain features of different types such as the altitude, the true airspeed or the engine power, sampled every second. To exploit these data, we propose to model statistically features of interest and to fit these models on the recorded data with off-the-shelf machine learning algorithms. The resulting estimators are then expected to take into account the actual performance of the aircraft, thus leading to a far more precise description of its performance.

To illustrate our method, we propose to model the drag and lift coefficients. It is common in the aeronautic literature to model the drag and lift forces through these coefficients, which quantify drag and lift independently from the wing size, airspeed, and air density (McCormick, 1995, chapter 2). In particular, the coefficients are used to somewhat capture very complex phenomena such as friction and permit to deduce the lift-to-drag ratio (Loftin, 1985, chapter 7), which plays a key role in assessing the performance of the aircraft.

The choice of these two aerodynamic coefficients is also motivated by the fact that neither of those coefficients is recorded by the aircraft (as a matter of fact, the drag and lift forces are also not recorded), highlighting the flexibility of this statistical approach. Indeed, this issue is bypassed by leveraging physical relationships to obtain approximated but explicit and deterministic formulas for the drag and lift coefficients. The statistical models are then fitted to approximated train data and their learning errors are computed on test sets. It is noteworthy that the models have to depend on features set by the user since they are aimed at being exploited by aeronautic softwares, such as those embedded in the Flight Management System of the aircraft. This explains in particular why physical models which may depend on other features, may not be used in such a setting.

This approach induces an additional error which we refer to as a physical approximation error, coming from the approximated data. To assess the prediction accuracy of the fitted models, this approximation error has to be taken into consideration. In a general setting within which we introduce our approach, we

propose bounds for the mean absolute error and relative error between the true value of the output and the predicted value from the model. These bounds depend explicitly on the physical approximation and the learning errors and are then applied in the present aeronautic setting. Note that in a slightly different setting, the problem can be interpreted as with errors on the response variable (Buonaccorsi, 1996), which is a particular case of the general framework of errors-in-variables models (Schennach, 2016; Fuller, 2009). For such problems, a statistical model is assumed on the distribution of the observed surrogate response variable given the unobserved response variable, for instance, the additive measurement error model (Carroll and Ruppert, 1988).

A similar approach to model unobserved aerodynamic variables has been proposed in Sun et al. (2018). In this paper, the authors aim at estimating the drag polar (i.e., a specific quadratic model for the drag coefficient depending on the lift one) by using a stochastic total energy model. Their approach is based on a Markov Chain Monte Carlo (MCMC) sampler to estimate posterior probability distributions of their parameters of interest. Similarly to our methodology, they exploit physical formulas to obtain approximate values for unobserved variables. Nevertheless, neither the associated error nor its impact on the prediction accuracy are taken into account in their analysis.

At this stage, let us stress that our aim is to propose a flexible methodology which is sufficiently generic to be used in any data-intensive engineering discipline. We show that this methodology can be straightforwardly applied to an aeronautic setting and that it leads to accurate predictions for two aerodynamic features, verifying in particular expected tendencies. This is illustrated with extensive numerical tests on recorded flight data. However, let us note that comparing the performance resulting from our estimated models to those provided by the manufacturers is unfortunately out of reach: in general, the manufacturers performance models are not publicly available (for commercial reasons) and therefore cannot be used for academic research.

The paper is organized as follows: in Section 2, we first propose an abstract formulation of the problem of modeling a variable for which only approximated data are available. Lemmas 1 and 2 provide the above mentioned bounds for the prediction error. Section 3 aims at specifying the aeronautic setting of interest. Let us mention that we restrict our study to the cruise phase for which physical approximation errors values are available, however, our method is actually not limited to this particular phase. Numerical results based on real flight data are presented and discussed in Section 4, and the paper closes on avenues for future work in Section 5.

2. Statistical Modeling

In this section, we consider a general setting where one aims at explaining a real-valued random variable $Y^* \in \mathcal{Y} \subseteq \mathbb{R}$ through a function f^* depending on the vector $X \in \mathcal{X} \subseteq \mathbb{R}^{d_X}$. We formulate this as the following regression problem:

$$Y^* = f^*(X) + \varepsilon, \quad (1)$$

where ε denotes a noise variable standing for unexplained determinants of Y^* . Nevertheless, in our setting, Y^* is a latent variable: no direct observation for Y^* is available, turning the direct estimation of f^* impossible. Our idea here is to propose an estimator for a surrogate of Y^* for which data can be obtained. To do so, suppose that there exists a relationship between Y^* and observed variables contained in a vector $Z \in \mathcal{Z} \subseteq \mathbb{R}^{d_Z}$, which can be observed together with X . More precisely, we suppose that there exists a known and explicit function $\varphi : \mathcal{Z} \rightarrow \mathbb{R}$ such that

$$\mathbb{E}_{Y^*, Z}[|Y^* - \varphi(Z)|] \leq r,$$

where $r > 0$ is known, and we let \mathbb{E}_A denote the expectation with respect to a random variable A . Thus, the variable $Y := \varphi(Z)$ can be considered as an approximation of Y^* , coming from a physical formula for instance. The error $\eta : \mathcal{Y}^* \times \mathcal{Z} \rightarrow \mathbb{R}$ of this approximation is defined as follows

$$\forall (y^*, z) \in \mathcal{Y}^* \times \mathcal{Z}, \quad \eta(y^*, z) := y^* - \varphi(z),$$

and will be named the *physical approximation error*. We consider then the following regression problem

$$Y = f(X) + \epsilon. \tag{2}$$

Note that $\varphi(Z)$ is actually a model for the latent variable Y^* , but our aim is to model this variable via the input X and not Z , motivating the problem (2). For instance, such a requirement arises in industrial contexts where inputs of some specialist softwares may differ from the variables involved in physical equations. If we assume that we have access to n random observations (x_i, z_i) (realizations from X and Z), we can derive observations for Y as follows

$$\forall i = 1, \dots, n, \quad y_i := \varphi(z_i),$$

leading to a training set $\mathcal{D} := (x_i, y_i)_{i=1}^n$. Contrary to the original problem (1) for which no training set is available, an estimator \hat{f} for the model f can be derived by solving the following minimization problem:

$$\hat{f} \in \arg \min_{g \in \mathcal{H}} \sum_{i=1}^n \ell(y_i, g(x_i)),$$

where the hypothesis class \mathcal{H} and the loss function $\ell : \mathbb{R}^2 \rightarrow \mathbb{R}$ are generic at this stage. For instance, one may consider the class of polynomials and the squared error loss.

Let us now upper bound the mean of the absolute value of the *total error*, defined by $Y^* - \hat{f}(X)$. In other words, the total error is the error between the unobserved variable Y^* and the predicted value $\hat{f}(X)$ given the training set \mathcal{D} . Note that the total error can be decomposed as follows:

$$Y^* - \hat{f}(X) = \eta(Y^*, Z) + (Y - \hat{f}(X)). \tag{3}$$

This is actually given by the sum of the physical approximation error $\eta(Y^*, Z)$ and another error term $Y - \hat{f}(X)$ which will be named the *learning error*. Indeed it comes from the statistical approximation of Y by $\hat{f}(X)$ and depends specifically on the training set \mathcal{D} , on the chosen model f and the algorithm to compute the estimator.

Lemma 1. *We have* $\mathbb{E}_{X, Y^*} [|Y^* - \hat{f}(X)|] \leq r + \mathbb{E}_{X, Z} [|Y - \hat{f}(X)|].$ (4)

Proof. By conditioning on Z , we have

$$\begin{aligned} \mathbb{E}_{X, Y^*} [|Y^* - \hat{f}(X)|] &= \mathbb{E}_Z [\mathbb{E}_{X, Y^* | Z} [|Y^* - \hat{f}(X)| | Z]] \\ &\leq \mathbb{E}_Z [\mathbb{E}_{X, Y^* | Z} [|\eta(Y^*, Z)| | Z]] + \mathbb{E}_Z [\mathbb{E}_{X, Y^* | Z} [|Y - \hat{f}(X)| | Z]] \\ &= \mathbb{E}_{Y^*, Z} [|\eta(Y^*, Z)|] + \mathbb{E}_{X, Z} [|Y - \hat{f}(X)|] \\ &\leq r + \mathbb{E}_{X, Z} [|Y - \hat{f}(X)|]; \end{aligned} \tag{5}$$

note that we have used the triangle inequality applied to (3) to obtain (5).

In the case where the order of magnitude of the learning error is smaller than the one of r , [Lemma 1](#) shows in particular that trying to compute a more precise estimator will have little consequence on the above bound of the total error.

We end this section by comparing the total error with the mean value of Y^* in the following lemma. More precisely, we upper bound the ratio between the means of the absolute value of the total error and of Y^* by an explicit and calculable quantity. This ratio, which can be reported as a percentage by multiplying it by 100, provides a relative measure of accuracy for the estimator \hat{f} . We also mention that it agrees with the Weighted Absolute Percentage Error (WAPE) in the classical case where \hat{f} is an estimator for Y^* .

Lemma 2. Suppose that $\mathbb{E}_Z[\varphi(Z)] > r$. Then $\mathbb{E}[Y^*]$ is positive and we have

$$\frac{\mathbb{E}_{X,Y^*}[|Y^* - \hat{f}(X)|]}{\mathbb{E}[Y^*]} \leq \frac{r + \mathbb{E}_{X,Z}[|Y - \hat{f}(X)|]}{\mathbb{E}_Z[\varphi(Z)] - r}. \quad (6)$$

Proof. We have

$$\mathbb{E}_{Y^*,Z}[\varphi(Z) - Y^*] \leq |\mathbb{E}_{Y^*,Z}[Y^* - \varphi(Z)]| \leq \mathbb{E}_{Y^*,Z}[|Y^* - \varphi(Z)|],$$

where we have applied Jensen's inequality to the absolute value function to obtain the second inequality. Moreover the linearity of the expected value and the assumption

$$\mathbb{E}_{Y^*,Z}[|Y^* - \varphi(Z)|] \leq r$$

lead to

$$\mathbb{E}_Z[\varphi(Z)] - r \leq \mathbb{E}[Y^*]. \quad (7)$$

Since $\mathbb{E}_Z[\varphi(Z)]$ is supposed to be larger than r , we deduce that $\mathbb{E}[Y^*] > 0$. Then we can take the inverse of inequality (7) and combine the result with inequality (4) to obtain (6). \square

In the following sections, we apply this abstract approach to model aerodynamic variables together with total error bounds. However it is noteworthy that this data-centric approach is sufficiently generic to be exploited in other disciplines.

3. Application to Aircraft Performance

We now move to modeling, the drag coefficient C_D^* and the lift coefficient C_L^* for a given narrow-body aircraft type for cruise conditions by exploiting recorded flight data¹. The predicted values of these coefficients are then expected to reflect real flights conditions. For instance, the coefficients C_D^* and C_L^* are used to establish the drag polar, which contains the aerodynamics of the aircraft (Anderson, 1999, section 2.9). As it is classically assumed in the aeronautics literature (see for instance Sun et al., 2018), our models for these coefficients will depend on the angle of attack α and on the Mach number M .

Nevertheless, the coefficients C_D^* and C_L^* are neither observed nor measured by the sensors of the aircraft during the flight. We therefore leverage the approach developed in Section 2. Following this approach, the main task is to determine approximated yet accurate formulas for the coefficients together with bounds for the physical approximation errors. With these approximations, we will be able to build data sets for approximated C_D^* and C_L^* which are expected to reflect on the actual aerodynamics of the aircraft. Models for C_D^* and C_L^* (depending on the angle of attack and Mach number) will then be trained on these datasets and their total errors will be bounded by using Lemma 1.

Prior to this, we emphasize that the method proposed in this paper is not limited to the present setting. It can be extended to other variables, aircraft types or phases, subject to available physical formulas and data.

For the sake of readability, Table 1 provides the names, the symbols and the SI units of the main physical variables used in the rest of this paper. In addition, the correspondence between the abstract variables and maps defined in Section 2 and the aeronautic ones is presented in Table 2.

¹ To be consistent with the notations introduced in Section 2, we let Y^* denote an exact but unobserved variable and we define the observed variable $Y := Y^* + \eta$, where η denotes an error term.

Table 1. Names, symbols and units of variables.

Variable name	Symbol	Unit (SI)
Angle of attack	α	rad
Path angle	γ	rad
True airspeed	V	m s^{-1}
Mach number	M	1
Altitude	h	m
Mass	m	kg
Fuel flow	FF	kg s^{-1}
Static air temperature	SAT	K
Air density	ρ	kg.m^{-3}
Thrust force	T	N
Drag force	D	N
Lift force	L	N

Table 2. Correspondence between abstract variables defined in Section 2 and the physical variables.

Abstract setting	Aeronautic setting
Y	C_D or C_L
X	(α, M)
Z	$(\rho, V, \alpha, \text{FF}, \text{SAT}, h, M, m, \gamma)$
φ	See (14) and (16)

Let us stress that the approximations we exploit here are actually derived from flight dynamics equations, whose accuracy depends on still existing physical models. In particular, we will use substantially the following approximated formula for the specific fuel consumption, noted here C_{SR}^* , from Roux (2005, p. 41):

$$C_{SR} := ((a_1(h)\lambda + a_2(h))M + (b_1(h)\lambda + b_2(h)))\sqrt{\frac{\text{SAT}}{\text{SAT}_0}} + (7.4e - 13(\varepsilon_c - 30)h + c)(\varepsilon_c - 30), \quad (8)$$

where

- SAT_0 is the temperature at sea level. Following the International Standard Atmosphere, it is set to 288.15 K;
- λ is the bypass ratio which depends on the turbofan engines; here this value is fixed because we consider a single airliner type;
- ε_c is the engine pressure ratio, which is also fixed here; and
- a_1, a_2, b_1, b_2 are linear piecewise functions (depending on the altitude) and c a constant which are given in Roux (2005, table 2.8).

As pointed out by Roux (2002), this model improves the classical one of Torenbeek (1982) and its mean relative error and its standard deviation for cruise conditions are given in Roux (2002, p. 66): they are equal respectively to 3.68% and 4.48%. Thus the coefficient C_{SR}^* satisfies the following equation

$$C_{SR}^* = C_{SR}(\text{SAT}, h, M) + \eta(C_{SR}^*, \text{SAT}, h, M) \quad (9)$$

where

$$\mathbb{E}_{C_{SR}^*, SAT, h, M} \left[\frac{|\eta(C_{SR}^*, SAT, h, M)|}{C_{SR}^*} \right] = 3.68 \times 10^{-2}, \tag{10}$$

over the cruise domain.

We establish now physical approximations for C_D^* and C_L^* in the case of a flight in a vertical plane and under the approximation that the Earth is locally flat. By applying Newton’s second law to a body (modeling the aircraft) of mass m moving in an air mass with no wind variations and by projecting the resulting equation onto the body frame, one obtains the following differential equations:

$$\begin{cases} m \dot{V} = T \cos \alpha - D - mg \sin \gamma \\ m V \dot{\gamma} = T \sin \alpha + L - mg \cos \gamma \end{cases}, \tag{11}$$

where g is the value of gravitational acceleration on Earth (here rounded to 9.81 m s^{-2}) and \dot{x} denotes the time-derivative of any physical variable x . We refer for instance to Rommel (2018) for a detailed derivation of the above relations. Moreover, we have the following relations:

$$\begin{cases} FF = C_{SR}^* T \\ D = \frac{1}{2} \rho V^2 S C_D^* \\ L = \frac{1}{2} \rho V^2 S C_L^* \end{cases}, \tag{12}$$

with S denoting the wing-surface of the aircraft; note that this value is fixed in our setting. From the system (12), we clearly have

$$\begin{cases} T = \frac{FF}{C_{SR}^*} \\ C_D^* = \frac{2}{\rho V^2 S} D \\ C_L^* = \frac{2}{\rho V^2 S} L \end{cases}.$$

Combining the system (11) with the preceding relations gives

$$\begin{cases} C_D^* = \frac{2}{\rho V^2 S} \left(\cos \alpha \frac{FF}{C_{SR}^*} - m \dot{V} - mg \sin \gamma \right) \\ C_L^* = \frac{2}{\rho V^2 S} \left(-\sin \alpha \frac{FF}{C_{SR}^*} + m V \dot{\gamma} + mg \cos \gamma \right) \end{cases}. \tag{13}$$

Apart from the specific fuel consumption C_{SR}^* , all the variables appearing in the right-hand sides of the system (13) are either recorded by the aircraft or easily calculable from other recorded variables via well-known physical relations.

By inserting Equation (9) into Equation (13), we obtain the following formulas for C_D^* and C_L^* :

$$\begin{cases} T = \frac{FF}{C_{SR}^*} - \frac{FF}{C_{SR}^*} \frac{\eta_{C_{SR}^*}}{C_{SR}^*} \\ C_D^* = \frac{2}{\rho V^2 S} \left(\cos \alpha \frac{FF}{C_{SR}^*} - m \dot{V} - mg \sin \gamma - \cos \alpha \frac{FF}{C_{SR}^*} \frac{\eta_{C_{SR}^*}}{C_{SR}^*} \right), \\ C_L^* = \frac{2}{\rho V^2 S} \left(-\sin \alpha \frac{FF}{C_{SR}^*} + m V \dot{\gamma} + mg \cos \gamma + \sin \alpha \frac{FF}{C_{SR}^*} \frac{\eta_{C_{SR}^*}}{C_{SR}^*} \right) \end{cases}$$

where $\eta_{C_{SR}^*} := \eta(C_{SR}^*, (SAT, h, M))$ for the sake of simplicity. By defining

- $Z := (\rho, V, \alpha, FF, SAT, h, M, m, \gamma)$;
- $\varphi_{C_D^*}(Z) := \frac{2}{\rho V^2 S} \left(\cos \alpha \frac{FF}{C_{SR}(SAT, h, M)} - m \dot{V} - mg \sin \gamma \right)$;
- $\eta_{C_D^*}(C_D^*, Z) := -\frac{2 \cos \alpha}{\rho V^2 S} \frac{FF}{C_{SR}(SAT, h, M)} \frac{\eta_{C_{SR}^*}}{C_{SR}^*}$,

(14)

we can write

$$C_D^* = \varphi_{C_D^*}(Z) + \eta_{C_D^*}(C_D^*, Z),$$

the variable $C_D := \varphi_{C_D^*}(Z)$ being the desired approximation for C_D^* . Similarly we obtain

$$C_L^* = \varphi_{C_L^*}(Z) + \eta_{C_L^*}(C_L^*, Z),$$
(15)

with

- $\varphi_{C_L^*}(Z) := \frac{2}{\rho V^2 S} \left(-\sin \alpha \frac{FF}{C_{SR}} + m V \dot{\gamma} + mg \cos \gamma \right)$;
- $\eta_{C_L^*}(C_L^*, Z) := \frac{2 \sin \alpha}{\rho V^2 S} \frac{FF}{C_{SR}(SAT, h, M)} \frac{\eta_{C_{SR}^*}}{C_{SR}^*}$.

(16)

Here the variable C_L^* is approximated by $C_L := \varphi_{C_L^*}(Z)$.

For the sake of readability, we sum up in [Figure 1](#) the relationships between the variables involved in the computations of C_D and C_L .

We now provide bounds for the means over the cruise phase of the absolute values of the physical approximation errors $\eta_{C_D^*}$ and $\eta_{C_L^*}$. Noting that these two variables are defined by a product in our setting, we can apply Hölder’s inequality (with a choice of exponents 1 and $+\infty$) to obtain

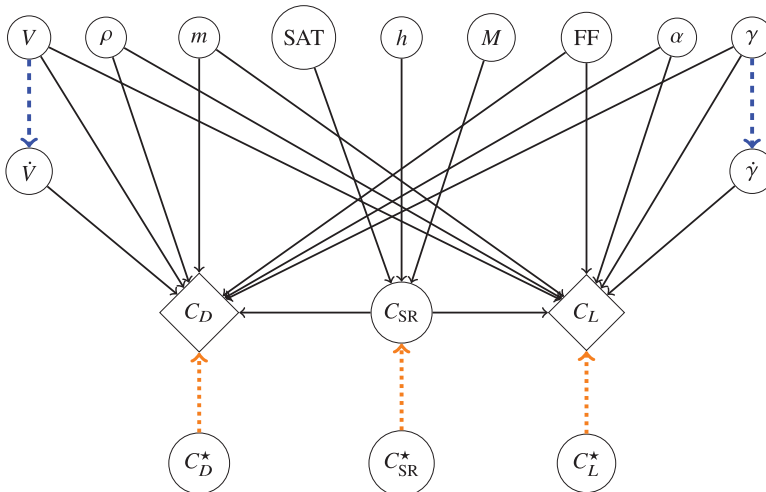


Figure 1. Relations between involved variables—black arrows correspond to deterministic relations, differentiation with respect to time is represented by blue dashed arrows and the orange dotted ones refer to physical approximations; variables in diamond-shaped boxes are the targets we aim at modeling.

$$\begin{cases} \mathbb{E}_{C_D^*, Z} \left[\left| \eta_{C_D^*}(C_D^*, Z) \right| \right] \leq K_{C_D^*} \mathbb{E}_{C_{SR}^*, SAT, h, M} \left[\frac{|\eta(C_{SR}^*, SAT, h, M)|}{C_{SR}^*} \right] \\ \mathbb{E}_{C_L^*, Z} \left[\left| \eta_{C_L^*}(C_D^*, Z) \right| \right] \leq K_{C_L^*} \mathbb{E}_{C_{SR}^*, SAT, h, M} \left[\frac{|\eta(C_{SR}^*, SAT, h, M)|}{C_{SR}^*} \right] \end{cases}, \tag{17}$$

where

$$K_{C_D^*} := \sup_{(\rho, V, \alpha, FF, SAT, h, M, m)} \left| \frac{2 \cos \alpha}{\rho V^2 S} \frac{FF}{C_{SR}(SAT, h, M)} \right|; \tag{18}$$

$$K_{C_L^*} := \sup_{(\rho, V, \alpha, FF, SAT, h, M, m)} \left| \frac{2 \sin \alpha}{\rho V^2 S} \frac{FF}{C_{SR}(SAT, h, M)} \right|. \tag{19}$$

The supremum is over the cruise domain here and the mean absolute relative error for C_{SR}^* is equal to $3.68e - 2$ according to (10). We refer to Section 4 for values of $K_{C_D^*}$ and $K_{C_L^*}$ computed from the available data.

Table 2 Similarly to Sun et al. (2018), it is possible to obtain simpler approximated formulas for C_D^* and C_L^* by assuming the following steady flight conditions with constant speed:

- 1.the altitude h is constant and so the path angle γ is equal to 0;
- 2.the angle of attack α is neglected: it is supposed to be equal to 0; and
- 3.the true airspeed V is constant and so its time-derivative \dot{V} is equal to 0.

In this case, we have:

$$\begin{cases} T = \frac{FF}{C_{SR}} - \frac{FF}{C_{SR}} \frac{\eta_{C_{SR}^*}}{C_{SR}^*} \\ C_D^* = \frac{2}{\rho V^2 S} \left(\frac{FF}{C_{SR}} - \frac{FF}{C_{SR}} \frac{\eta_{C_{SR}^*}}{C_{SR}^*} \right) \\ C_L^* = \frac{2mg}{\rho V^2 S} \end{cases}.$$

In the present paper, we consider the complete formulas (13) which are likely to best preserve accuracy of the approximations and to catch real flight conditions.

4. Experiments

In this section, we present numerical results based on real flight data for the method introduced in Section 2 and applied to the aeronautic setting described in Section 3. We first detail the data and the preprocessing steps we carried out, before reporting experiments design and results.

4.1. Data description and preprocessing

We have access to 423 recorded short and medium-haul flights performed by the same narrow-body airliner, the data being recorded by the QAR. These flight data are provided by a partner airline and cannot be publicly released for commercial reasons. From this dataset, we extract all the observations for the variables contained in the vector Z defined in Section 3. The heading and the wind speed are also extracted to remove heading changes and high wind variations. All these variables are then smoothed by means of smoothing splines to remove the noise coming from measuring instruments and converted into the international system of units. The time-derivatives are computed on the basis of the smoothing splines. As

explained in Section 3, we consider cruise phases in a vertical plane with no wind variations, so we require the following conditions to be satisfied:

- we keep observations from the top of climb to the top of descent without those corresponding to climb steps; from a numerical point of view, we keep time-intervals such that the standard deviations of the altitude over these intervals is smaller than an arbitrary small threshold;
- the heading angle of the aircraft has to be constant; from a numerical point of view, we keep intervals such that the standard deviations of the heading over these intervals is smaller than an arbitrary small threshold;
- the wind speed variations have to be equal to 0; from a numerical point of view, we keep intervals such that the means and the standard deviations of the time-derivative of the wind over these intervals are smaller than an arbitrary small threshold; and
- the lengths of the resulting intervals have to be larger than 10 s.

Given the time-intervals during which the above conditions are satisfied, we sample every 10 s in each interval. This is motivated by the fact that the errors of the resulting models trained on the dataset sampled every 10 s and on the dataset without sampling are very close. Hence, sampling allows to reduce the learning time without impacting strongly the accuracy. Afterwards, the values for the approximated variables C_D and C_L are computed by means of the functions $\varphi_{C_D^*}$ and $\varphi_{C_L^*}$ defined in Section 3. Finally, we have 164,054 observations which are randomly split into training, validation, and test sets (70% of the dataset is used for the training, 20% for the validation, and 10% for the test).

Table 3 presents an example of a preprocessed data set (with simulated values to avoid divulging the dataset).

4.2. Experiments design

Here, we aim at estimating the following models for the approximated drag C_D and lift C_L coefficients

$$C_D = f_{C_D}(\alpha, M) \quad \text{and} \quad C_L = f_{C_L}(\alpha, M) ,$$

by exploiting the preprocessed data described above. To do so, we consider different classical models which are introduced in Table 4. This table also gives the considered hyper-parameters and their range. The hyper-parameters are tuned by using threefold cross-validation, the loss function being the mean squared error. Furthermore, we use an early stopping rule when fitting the gradient tree boosting model to limit the number of iterations, the validation set being used to stop iterating. The maximum number of iterations has been set to 5,000 in this case. In the end, we are interested in the three following classical error metrics: the root-square of the mean squared error, the mean absolute error and the mean absolute percent error. We use the software package `LightGBM` (Ke et al., 2017) as an implementation of the gradient tree boosting algorithm and we compute the other models using `scikit-learn` Python library (Pedregosa et al., 2011).

Table 3. Example of a preprocessed dataset.

Observation	ρ	V	α	FF	...	m	γ
1	0.3224	234.5	0.0324	0.6716	...	62,519	0.0139
2	0.3704	236.8	0.0224	0.6503	...	64,960	0.0198
3	0.3224	234.8	0.0305	0.6637	...	66,974	0.0159
⋮	⋮	⋮	⋮	⋮		⋮	⋮
164,054	0.3433	232.9	0.0332	0.6642	...	66,673	0.0150

Table 4. Hyper-parameters and their range for the considered models.

Model	Hyper-parameters	Range
Constant	None	\emptyset
Linear	None	\emptyset
Polynomial	Degree	{2, 3, 4, 5}
SVM	Kernel	{linear, polynomial, Gaussian, sigmoid}
k-NN	Neighbors number, weights	{1, 21, 41, ..., 701} \times {uniform, distance}
Decision tree	Trees depth	{1, 2, ..., 10}
Random forest	Trees depth, trees number	{1, 2, ..., 6} \times {100, 200, ..., 700}
Gradient tree boosting	Trees depth	{1, 2, ..., 6}

Table 5. Means and standard deviations of error metrics for different C_D models computed over 100 independent repetitions.

C_D model	RMSE	MAE	MAPE [%]
Constant	$(8.778 \pm 0.285) \times 10^{-3}$	$(5.932 \pm 0.102) \times 10^{-3}$	53.58 ± 56.58
Linear	$(1.992 \pm \mathbf{0.021}) \times 10^{-3}$	$(1.424 \pm \mathbf{0.006}) \times 10^{-3}$	$4.53 \pm \mathbf{0.05}$
Polynomial	$(\mathbf{1.926} \pm 0.029) \times 10^{-3}$	$(1.361 \pm 0.008) \times 10^{-3}$	4.31 ± 0.06
SVM	$(1.392 \pm 0.259) \times 10^{-2}$	$(1.360 \pm 0.269) \times 10^{-2}$	43.25 ± 8.86
k-NN	$(1.936 \pm 0.025) \times 10^{-3}$	$(1.362 \pm \mathbf{0.006}) \times 10^{-3}$	$4.31 \pm \mathbf{0.05}$
Decision tree	$(1.961 \pm 0.038) \times 10^{-3}$	$(1.364 \pm 0.009) \times 10^{-3}$	$4.32 \pm \mathbf{0.05}$
Random forest	$(1.929 \pm 0.026) \times 10^{-3}$	$(1.357 \pm 0.007) \times 10^{-3}$	$\mathbf{4.29} \pm \mathbf{0.05}$
Gradient tree boosting	$(1.928 \pm 0.047) \times 10^{-3}$	$(\mathbf{1.356} \pm 0.012) \times 10^{-3}$	$\mathbf{4.29} \pm 0.08$

The smallest values are indicated by bolded numbers.

Abbreviations: MAE, mean absolute error; MAPE, mean absolute percent error; RMSE, mean squared error.

4.3. Results

We performed 100 times the learning process: at each time, the preprocessed data is randomly split into training, validation and test sets and the models are estimated and tested using these sets. The means and the standard deviations of the errors computed on the test sets are given in Tables 5 and 6. We report numbers up to a precision of three decimal digits.

Figures 2–4 allow to visualize the tendencies of estimators \hat{f}_{C_D} and \hat{f}_{C_L} with respect to the Mach number for different fixed values of the angle of attack. Figures 2–4 show respectively a polynomial, a decision tree and a gradient tree boosting models.

First of all, we observe that the decision tree and gradient tree boosting models lead to raw predicted curves which may be hard to interpret from an aeronautic point of view. This is especially true for the decision tree model even though its learning error is similar to those of the two other models. Since we aim here at checking whether some expected aeronautic tendencies are caught by our approach, we smooth the predicted curves by means of smoothing splines to interpret the results in an easier way. These smoothed curves are given by the dotted curves in Figures 3 and 4. Note that the use of smoothing splines is to highlight the underlying tendency, but is certainly not used as a predictive method because of well-known artefacts due to the smoothing technique.

Now we mention that 90% of Mach number data are between 0.77% and 0.80% and 90% of angle of attack data are between 1.9° and 2.9°. Then we observe that both predicted C_D and C_L globally increase

Table 6. Means and standard deviations of error metrics for different C_L models computed over 100 independent repetitions.

C_L model	RMSE	MAE	MAPE [%]
Constant	$(7.358 \pm 0.116) \times 10^{-2}$	$(5.999 \pm 0.060) \times 10^{-2}$	14.24 ± 0.14
Linear	$(1.436 \pm 0.007) \times 10^{-2}$	$(1.102 \pm 0.004) \times 10^{-2}$	$2.17 \pm \mathbf{0.01}$
Polynomial	$(1.205 \pm \mathbf{0.005}) \times 10^{-2}$	$(9.203 \pm \mathbf{0.036}) \times 10^{-3}$	$1.78 \pm \mathbf{0.01}$
SVM	$(6.272 \pm 0.029) \times 10^{-2}$	$(5.982 \pm 0.030) \times 10^{-2}$	11.16 ± 0.05
k-NN	$(\mathbf{1.181} \pm 0.006) \times 10^{-2}$	$(8.976 \pm 0.039) \times 10^{-3}$	$1.73 \pm \mathbf{0.01}$
Decision tree	$(1.185 \pm 0.008) \times 10^{-2}$	$(\mathbf{8.886} \pm 0.045) \times 10^{-3}$	$\mathbf{1.71} \pm \mathbf{0.01}$
Random forest	$(1.209 \pm 0.006) \times 10^{-2}$	$(9.211 \pm 0.041) \times 10^{-3}$	$1.78 \pm \mathbf{0.01}$
Gradient tree boosting	$(1.192 \pm 0.011) \times 10^{-2}$	$(9.025 \pm 0.072) \times 10^{-3}$	1.75 ± 0.02

The smallest values are indicated by bolded numbers.

Abbreviations: MAE, mean absolute error; MAPE, mean absolute percent error; RMSE, mean squared error.

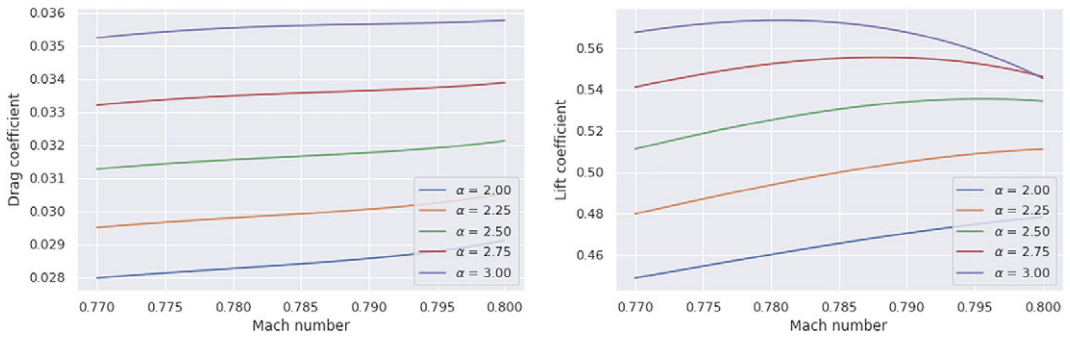


Figure 2. Predictions of C_D and C_L from polynomial models.




Figure 3. Predictions of C_D and C_L from decision trees models. Solid lines are the raw prediction curves and dotted lines are smoothed versions (using splines).

when the Mach number or the angle of attack increases. This global tendency is actually expected in this small range of values according to Anderson (1999, part 1, chapter 2): the larger the angle of attack or the Mach number, the larger the drag and lift coefficients.

Nevertheless, this natural tendency for the lift coefficient is not verified by the estimators when α is too large, namely $\alpha = 2.75^\circ$ or $\alpha = 3^\circ$. This unexpected behavior can be explained by the approximated nature




Figure 4. Predictions of C_D and C_L from decision gradient tree boosting models. Solid lines are the raw prediction curves and dotted lines are smoothed versions (using splines).

of the variable C_L . Indeed it may behave in a way that is different from C_L^* in certain regions of the cruise domain. In this case, any estimator for C_L is likely to inherit this unexpected behavior and we believe refined aeronautics-supported approximations would bring a solution.

We now focus on the physical approximation errors for the drag and lift coefficients, that is to say $\eta_{C_D^*}$ and $\eta_{C_L^*}$. According to (17), the mean of the absolute value of these errors is bounded by the product between the constants $K_{C_D^*}$ or $K_{C_L^*}$ with the mean absolute relative error of C_{SR}^* , the latter being equal to 3.68×10^{-2} . We estimate the value of these constants by using our recorded observations: the maximal values of $K_{C_D^*}$ or $K_{C_L^*}$ (defined in Equations 18 and 19) are respectively equal to 4.38×10^{-2} and 2.94×10^{-3} . Using these two values as estimators for $K_{C_D^*}$ and $K_{C_L^*}$ gives the following bounds for the physical approximation errors:

$$\begin{cases} \mathbb{E}_{C_D^*, Z} [|\eta_{C_D^*}(C_D^*, Z)|] \leq 1.61 \times 10^{-3} \\ \mathbb{E}_{C_L^*, Z} [|\eta_{C_L^*}(C_L^*, Z)|] \leq 1.08 \times 10^{-4} \end{cases}$$

According to the generic inequality given in Lemma 1, a bound for the mean of the absolute total error (defined in Section 2) of a given variable can be obtained by adding up the bounds for the physical approximation and learning errors. The latter is here approximated by the MAE of the estimated model. To provide an example of numerical bounds for the total errors of the drag and lift coefficients, we choose estimators \hat{f}_{C_D} and \hat{f}_{C_L} whose MAE values are equal to the MAE means given in Tables 5 and 6. Note that this choice is motivated by the fact that the standard deviations of the MAE for the different models are much smaller than the means. Following this choice, examples of bounds for the total errors are then given in Table 7.

The empirical means of $C_D = \varphi_{C_D^*}(Z)$ and $C_L = \varphi_{C_L^*}(Z)$ over our data are respectively equal to

$$\widehat{\mathbb{E}}_Z [\varphi_{C_D^*}(Z)] = 3.23 \times 10^{-2}, \quad \widehat{\mathbb{E}}_Z [\varphi_{C_L^*}(Z)] = 5.32 \times 10^{-1},$$

showing in particular that $\widehat{\mathbb{E}}_Z [\varphi_{C_D^*}(Z)] > 1.61 \times 10^{-3}$ and $\widehat{\mathbb{E}}_Z [\varphi_{C_L^*}(Z)] > 1.08 \times 10^{-4}$ (we used the slight notation abuse $\widehat{\mathbb{E}}$ to denote the empirical mean). The hypotheses of Lemma 2 are then satisfied and we are in position to upper bound the ratios between the mean of the absolute value of the total errors and the mean values of C_D^* and C_L^* . Numerical values expressed as a percentage for these bounds are given in Table 7.

These results capture how classical, off-the-shelf regression methods perform and somewhat surprisingly, most methods compete on similar grounds (to the notable exception of the SVM). This suggests that most of these methods achieve a good enough complexity to capture the underlying phenomenon.

To finish, we mention that our approach is sufficiently generic to be applied to other settings, such as other flight phases. For the sake of illustration, we consider the drag coefficient during the climb phase. In this case, we exploit the climb data from our 423 available recorded flights, that is to say data for which the altitude is between

Table 7. Bounds for the mean absolute and mean relative total errors of the drag and lift coefficients using estimators \hat{f}_{C_D} and \hat{f}_{C_L} whose mean absolute error (MAE) values are equal to the MAE means given in Tables 5 and 6—*Absolute* and *Relative* refer respectively to the bounds given in Lemmas 1 and 2.

Models	Drag coefficient		Lift coefficient	
	Absolute	Relative [%]	Absolute	Relative [%]
Constant	7.542×10^{-3}	24.57	6.010×10^{-2}	11.30
Linear	3.034×10^{-3}	9.89	1.113×10^{-2}	2.09
Polynomial	2.971×10^{-3}	9.68	9.311×10^{-3}	1.75
SVM	1.521×10^{-2}	49.56	5.993×10^{-2}	11.27
k-NN	2.972×10^{-3}	9.68	9.084×10^{-3}	1.71
Decision tree	2.974×10^{-3}	9.69	8.994×10^{-3}	1.69
Random forest	2.967×10^{-3}	9.67	9.319×10^{-3}	1.75
Gradient tree boosting	2.966×10^{-3}	9.67	9.133×10^{-3}	1.72

Table 8. Means and standard deviations of error metrics for different C_D models for climb phase computed over 100 independent repetitions.

C_D model	RMSE	MAE	MAPE [%]
Constant	$(8.912 \pm 0.373) \times 10^{-3}$	$(5.919 \pm 0.102) \times 10^{-3}$	48.21 ± 40.71
Linear	$(6.061 \pm 0.208) \times 10^{-3}$	$(3.618 \pm 0.041) \times 10^{-3}$	$29.08 \pm \mathbf{14.16}$
Polynomial	$(5.611 \pm 0.215) \times 10^{-3}$	$(3.399 \pm 0.037) \times 10^{-3}$	27.94 ± 14.33
SVM	$(2.274 \pm 0.708) \times 10^{-2}$	$(1.937 \pm 0.659) \times 10^{-2}$	82.99 ± 24.74
k-NN	$(5.639 \pm 0.215) \times 10^{-3}$	$(3.420 \pm \mathbf{0.035}) \times 10^{-3}$	26.90 ± 14.93
Decision tree	$(5.840 \pm 0.315) \times 10^{-3}$	$(3.494 \pm 0.050) \times 10^{-3}$	26.74 ± 14.26
Random forest	$(5.598 \pm 0.189) \times 10^{-3}$	$(3.428 \pm \mathbf{0.035}) \times 10^{-3}$	27.02 ± 16.70
Gradient tree boosting	$(5.678 \pm 0.425) \times 10^{-3}$	$(3.405 \pm 0.073) \times 10^{-3}$	33.27 ± 32.07

The smallest values are indicated by bolded numbers.
 Abbreviations: MAE, mean absolute error; MAPE, mean absolute percent error; RMSE, mean squared error.

3,000 ft and the top of climb, and we apply the same preprocessing and learning steps as those described in this section. The results for the means and standard deviations of the learning errors computed over a test set are given in Table 8. We remark that these errors are much larger than those for the cruise phase. This accuracy loss can be explained by the fact that each variable during the climb phase has a larger range. For instance, the Mach number varies from 0.3 at the beginning of the climb to 0.81 at the top of climb while it varies only from 0.76 and 0.81 during the cruise. In addition, we are not in position to compute a numerical value for a bound of the physical approximation errors means for C_D^* . This is due to the fact that Roux (2002) does not estimate the physical approximation error mean coming from its model for the variable C_{SR}^* for the climb phase. Once such a quantity is available, numerical bounds for the total errors for C_D^* can be derived in this case.

5. Conclusion

Our contributions are twofold: (a) we have proposed individual models trained on in-air data to improve the current aeronautic performance of individual aircrafts, rather than industry-wide calibrated parameters. This

allows in particular for the search of more efficient (e.g., flight duration, speed, fuel consumption, etc.) trajectories for aircrafts and (b) we have designed a generic framework combining off-the-shelf machine learning with domain-specific approximations, which can be used in any data-intensive engineering discipline. We certainly hope that this approach can be replicated in other fields of study. We also intend to use this approach as a building block to optimizing end-to-end pipelines, for example for in-air real-time fuel optimization.

Funding Statement. This work was supported by the PERF-AI project funded by the European research program Clean Sky 2. The funder had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interest. The authors declare no competing interests.

Data Availability Statement. Data that support the findings of this study have been gathered from the Quick Access Recorder (QAR) of an aircraft operated by a private airline, hence can not be released publicly for commercial reasons.

Authorship Contributions. Formal analysis, F.D., B.G., and V.V.; Methodology, F.D., B.G., and V.V.; Software, F.D.; Writing-original draft, F.D.; Writing-review & editing, B.G. and V.V. All authors approved the final submitted draft.

Acknowledgments. The authors are grateful to Baptiste Gregorutti and Pierre Jouniaux for fruitful discussions about the conceptualization of this research and the aeronautic validation of the results. The authors are grateful to Arthur Talpaert for his substantial contribution to the code used to implement our methods.

Supplementary Materials. To view supplementary material for this article, please visit <http://dx.doi.org/10.1017/dce.2020.12>.

References

- Airbus** (2001) *Getting Hands on Experience with Aerodynamic Deterioration*. Technical report. Blagnac, France: Airbus. http://www.smartcockpit.com/docs/Getting_Hands-On_Experience_With_Aerodynamic_Deteriorations.pdf.
- Airbus** (2002) *Getting to Grips with Aircraft Performance Monitoring*. Technical report. Blagnac, France: Airbus. http://www.smartcockpit.com/docs/Getting_to_Grips_With_Aircraft_Performance.pdf.
- Anderson J** (1999) *Aircraft Performance and Design*. McGraw-Hill International Editions: Aerospace Science/Technology Series. WCB/McGraw-Hill.
- Buonaccorsi JP** (1996) Measurement error in the response in the general linear model. *Journal of the American Statistical Association* 91(434), 633–642.
- Carroll RJ and Ruppert D** (1988) *Transformation and Weighting in Regression*, Vol. 30. CRC Press.
- Fuller WA** (2009) *Measurement Error Models*, Vol. 305. John Wiley & Sons.
- Ke G, Meng Q, Finley T, Wang T, Chen W, Ma W, Ye Q and Liu T-Y** (2017) Lightgbm: A highly efficient gradient boosting decision tree. In Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S and Garnett R (eds), *Advances in Neural Information Processing Systems*. Vol. 30. Curran Associates, Inc., pp. 3146–3154.
- Krajčák K, Nikolić D and Domitrović A** (2015) Aircraft performance monitoring from flight data. *Tehnički Vjesnik* 22(5), 1337–1344.
- Loftin LK** (1985) *Quest for Performance: The Evolution of Modern Aircraft*. NASA Scientific and Technical Information Branch.
- McCormick BW** (1995) *Aerodynamics, Aeronautics, and Flights Mechanics*, 2nd Edn. Wiley.
- Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M and Duchesnay E** (2011) Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12, 2825–2830.
- Rommel C** (2018) *Exploration de données pour l'optimisation de trajectoires aériennes*. PhD thesis. École Polytechnique.
- Roux E** (2002) *Modèles Moteurs... Réacteurs double flux civils et réacteurs militaires à faible taux de dilution avec Post-Combustion*. PhD thesis. INSA-SupAéro-ONÉRA.
- Roux E** (2005) *Pour une approche analytique de la Dynamique du Vol*. PhD thesis. SupAéro.
- Schennach SM** (2016) Recent advances in the measurement error literature. *Annual Review of Economics* 8(1), 341–377.
- Sun J, Hoekstra JM and Ellerbroek J** (2018) Aircraft drag polar estimation based on a stochastic hierarchical model. In *Eighth SESAR Innovation Days*.
- Torenbeek E** (1982) *Synthesis of Subsonic Airplane Design*. Delft University Press.

Cite this article: Dewez, F. Guedj, B. and Vandewalle, V. 2020. From industry-wide parameters to aircraft-centric on-flight inference: Improving aeronautics performance prediction with machine learning. *Data-Centric Engineering*, 1: e11. doi:<https://doi.org/10.1017/dce.2020.12>