



**POLITECNICO**  
MILANO 1863

[RE.PUBLIC@POLIMI](mailto:RE.PUBLIC@POLIMI)

Research Publications at Politecnico di Milano

## Post-Print

This is the accepted version of:

G. Gori, T. Davoli, A. Rausa, A. Zanotti, F. Auteri, A. Guardone  
*Bayesian Calibration of a Low Order Aerodynamic Model for the Design of Unconventional Tail Empennages*  
in: AIAA Aviation 2023 Forum, AIAA, 2023, ISBN: 9781624107047, p. 1-18, AIAA 2023-4422  
[AIAA Aviation 2023 Forum, San Diego, CA, USA, 12-16 June 2023]  
doi:10.2514/6.2023-4422

The final publication is available at <https://doi.org/10.2514/6.2023-4422>

Access to the published version may require subscription.

**When citing this work, cite the original published paper.**

Permanent link to this version

<http://hdl.handle.net/11311/1243979>

# Bayesian calibration of a low order aerodynamic model for the design of unconventional tail empennages

Giulio Gori\* and Andrea Rausa† and Alex Zanotti‡ and Franco Auteri§ and Alberto Guardone¶  
*Department of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, 20156, Milano, Italy.*

**In the frame of the MONNALISA Project of the Clean Sky 2 Joint Undertaking (CSJU), we develop a calibration framework for improving the prediction accuracy of a low-order aerodynamic model. The enhancement is sought with respect to a reference application case, an unconventional “Advanced Rear End” component for the forthcoming generation of ultra-efficient aircraft, which is a specific focus of the MONNALISA Project. The calibration framework develops upon a classical Bayesian inference approach providing not only optimal values for the model parameters, but also a probabilistic characterization in terms of their peculiar probability distribution within well defined bounds. The reference data include both numerical predictions from computational fluid dynamics simulations and experimental measurements collected at the large wind-tunnel facility (GVPM) at Politecnico di Milano.**

## I. Nomenclature

$Re$	=	Reynolds number
$M$	=	Mach number
$C_D$	=	Drag coefficient
$C_L$	=	Lift coefficient
$T_\infty$	=	Freestream temperature
$b$	=	Wing span
$\alpha$	=	Wing angle of attack
$\delta$	=	Wing control deflection
$\lambda$	=	Wing aspect ratio
$\tau$	=	Wing taper ratio
$\Gamma$	=	Wing dihedral angle
$\Lambda$	=	Wing sweep angle

## II. Introduction

THE mission of the Clean Sky 2 Joint Undertaking (CSJU) is to develop innovative technologies to improve the sustainability of commercial air mobility [1]. The main goal is to achieve a significant reduction of harmful pollutants released by the aviation industry. In particular, it is expected a cut of fuel burn and its related CO<sub>2</sub> emissions by 20-30%, reducing also the aeroacoustic footprint.

The MONNALISA Project [2] is funded in the frame of the CSJU and aims at developing a physics-based low-order model to allow the fast and accurate prediction of the nonlinear aerodynamic characteristics of innovative tailplane configurations for commercial aircraft. As such, the low-order model will contribute to the discovery of new concepts of rear end for advanced and ultra-advanced long-range and short/medium-range aircraft, possibly leading to improved aerodynamic performances of the future aircraft generation. In particular, the targeted application is a nonconventional tailplane configuration with a deflecting control surface at the trailing edge, and the low-order model considered here is expected to serve the optimal design process of such device.

\*Assistant Professor, Department of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, 20156, Milano, Italy.

†Ph.D. Student, Department of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, 20156, Milano, Italy.

‡Associate Professor, Department of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, 20156, Milano, Italy.

§Associate Professor, Department of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, 20156, Milano, Italy.

¶Full Professor, Department of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, 20156, Milano, Italy.

In order to achieve these goals, the MONNALISA project develops a multidisciplinary framework entailing three different disciplines, strictly entwined for the purpose of improving prediction tools:

- 1) **Experimental campaigns** devoted to collecting data through a systematic series of wind tunnel tests covering a wide range of tailplane configurations.
- 2) **Computational Fluid Dynamics (CFD)** simulations of unconventional tailplane configurations to increase the resolution of the experimental database.
- 3) **Bayesian calibration techniques** for enhancing the predictions of low-order models based on the available data.

This paper deals specifically with the third item of the list. Namely, it develops a computational framework for calibrating a low-order aerodynamic model, and it presents an assessment of the attained improvements.

In Sec. III we present the low-order aerodynamic model subject to calibration, which is implemented into a computer code. In Sec. IV we present the application test case, for which reference data (employed in the calibration process) are available from both experimental and high-fidelity CFD investigations. Section V presents the methodology underlying the calibration framework. Section VI reports a verification of the framework and the results achieved by virtue of the calibration of the low-order model, including a prediction comparison against the non-calibrated model.

### III. The aerodynamic low-order model

The aerodynamic low-order model employed in this work is implemented within the DUST solver. DUST is a flexible computational tool designed to obtain fast and reliable predictions concerning the aerodynamic performances of aircraft. It is an open source software<sup>\*</sup> resulting from a collaborative research effort between Politecnico di Milano and A<sup>3</sup> by Airbus LLC [3]. DUST particularly suits applications of great complexity because of its capability of accurately simulating arbitrary aircraft configurations, e.g., multi-rotor machines including moving parts.

Several aerodynamic models based on the potential flow assumption are available within the DUST solver, e.g., surface panels, Vortex Lattice Methods, and the Lifting Line method. In this paper, we take advantage of a modified Lifting Line (LL) approach which is briefly described in the following. According to the so-called  $\Gamma$ -method, the circulation of the lifting line is determined by solving a nonlinear problem requiring the knowledge of the aerodynamic coefficients characterizing each of the lifting sections. These aerodynamic coefficients are provided in the form of tabulated data expressing the dependency of the performance from the local angle of attack, the local Reynolds number, and the local Mach number. These tables can be either extracted from fully three-dimensional simulations of the cantilever wing, or from purely two-dimensional simulations of the sectional airfoil [4]. In this latter case, a correction must be implemented in order to retrieve three-dimensional flow effects. As mentioned, in this work we rely on a modified LL model implementing the sweep-angle correction proposed by Goitia and Llamas in [5]. The reader is referred to the original work for the details concerning the methodology. Here, we recall only the features necessary to illustrate the purpose of this work.

This correction extends the limits of applicability for the LL approach by improving the modeling of three-dimensional effects due to the sweep angle of the wing, which produces a spanwise boundary layer flow from the root to the tip if the sweep is positive. Because of the sweep angle  $\Lambda$ , the boundary layer is thinner at the root, introducing a delay of stall [6]. At the tip, the boundary layer is instead thicker causing an anticipation of the onset of stall. The classical sweep theory for infinite wings, see for instance [7], claims that purely 2D characteristics may be corrected and transformed into 2.5D ones through the relations

$$C_L^{2.5D} = C_L^{2D} \cos^2 \Lambda, \quad (1)$$

$$C_D^{2.5D} = C_D^{2D}, \quad (2)$$


$$\alpha^{2.5D} = (\alpha^{2D} - \alpha_0) \cos \Lambda + \alpha_0, \quad (3)$$

$$(4)$$

being  $C_L$  and  $C_D$  the lift and drag coefficients, respectively, and  $\alpha_0$  the zero-lift angle for the 2D sectional airfoil. The correction reported above is valid for infinite wings.

<sup>\*</sup><https://www.dust-project.org/>

<sup>†</sup>In our computational framework, these lookup tables are evaluated by means of SU2, using a two-dimensional RANS model of the sectional profile coupled to an in-house mesh adaptation software named Darwin, see [4]. The sequential mesh adaptation strategy is based on local error estimates of the pressure gradient.



**Fig. 1 Tail-plane model installed in the wind tunnel**

The correction suggested in [5] consists in implementing a further empirical corrective function for finite wings. In a finite swept wing, the aerodynamic characteristics of each section vary along the span coordinate  $y$  e.g., the maximum lift coefficient is larger at the root than at the tip [6]. Because of this, an additional correction  $\kappa = \kappa(y)$  is designed to vary along the spanwise direction. The two-dimensional polars are therefore assigned an effective sweep angle

$$\Lambda_{\text{eff}}(y) = \Lambda_{\text{loc}}(y) \cdot \kappa(y). \quad (5)$$

The results from high-accuracy CFD simulations [8] suggest implementing  $\kappa(y)$  as the superimposition of two different and independent contributions which become active at the root and at the tip

$$\kappa(y) = 1 + \kappa_{\text{root}}(\eta_{\text{root}}) + \kappa_{\text{tip}}(\eta_{\text{tip}}), \quad (6)$$

noting the coordinate transformation occurring through the definition of non-dimensional coordinates  $\eta_{\text{root}}$  and  $\eta_{\text{tip}}$ , see [5]. In the original paper, both the tip and the root functions are modeled using Bézier curves [9] defined as follows:

$$B(y) = C_1 (1 - y)^3 + 3C_2 (1 - y)^2 y + 3C_3 (1 - y) y^2 + C_4 y^3. \quad (7)$$

Therefore, the two curves are parametrized by means of four coefficients each. A variation of these parameters changes the shape of the spanwise correction accounting for three-dimensional effects associated to a finite swept wing. Summarizing, the considered aerodynamic low-order model i.e., the corrected LL model, entails a total of 8 free parameters which may be tuned to fit the available data at best. In the frame of this work, these parameters are indeed unknown and subject to calibration.

#### IV. The application case and available data

The application test case consists of a cantilever wing mounted vertically on the floor of the wind tunnel test section, see Fig. 1. In the frame of the MONNALISA Project, this wind-tunnel model underwent several investigations including both experiments carried out at the Politecnico di Milano Wind Tunnel Laboratory (GVPM) [10] and high-fidelity CFD analyses [11].

The parameters specifying the model geometry and the test conditions of the referenced works are summarized in Tab. 1. As clear from the table, the aerodynamic investigation targets discrete points entailing variable operative parameters namely,  $\alpha$ ,  $\delta$  and  $\Gamma$ , spanning specific ranges. In particular, only two dihedral angles are considered namely,  $\Gamma = 0^\circ$  and  $\Gamma = 45^\circ$ . At the same time, the  $\alpha$  range is investigated using a resolution of 2 angular degrees, whereas a resolution of 5 (in CFD) or 10 (in experiments) angular degrees is employed to span the range of the control surface deflection  $\delta$ . In our application case, the planform of the wing is fixed and equal for all the experiments (either physical or numerical). The geometry of the airfoil defining the section of the wing, as well as the full data set, is available under the ERCOFTAC community established on Zenodo [12]. Each folder contains the  $\alpha$  sweep referred to a single tail

**Table 1 Test, geometric and operative parameters for the reference test case.**

Test parameters			Geometric parameters			Operative parameters		
	Value	Dimensions		Value	Dimensions		Value	Dimensions
$Re$	$5.5 \times 10^6$	-	$\lambda$	5	-	$\alpha$	[0, 26]	deg
$M$	0.147	-	$\tau$	0.4	-	$\delta$	[-35, 35]	deg
$T_\infty$	300	K	$\Gamma$	[0, 45]	deg			
			$\Lambda$	30	deg			
			$b$	1.6	m			

geometry, defined by the geometric parameters and the deflection of the control surface. The data available include the flow solution in .vtu format, the computational mesh in CGNS format, CAD files in IGES, STEP and CATIAv5 formats, and a table containing the aerodynamic coefficients for each value of  $\alpha$ .

In the experimental campaign, different measurements has been carried out [10]. Namely, the wing integral forces and moments were measured by a six-component balance, while the hinge moment acting on the tail-plane control surface has been measured using a purposely built and calibrated strain-gauge torque meter. The boundary-layer transition line was instead exposed by means of an infrared thermography technique. Particle Image Velocimetry (PIV) surveys were also carried out for some selected test conditions with the movable surface oscillating. For the purposes of this paper, we take advantage of load measurements concerning the lift generated by the half-wing.

Besides the experimental campaign, this test case was also investigated numerically using a high-fidelity CFD solver based on the Reynolds-Averaged Navier-Stokes (RANS) model from the SU2 [13] open-source suite. An additional database including CFD results from 420 simulations, carried out at different points spanning the ranges reported in Tab 1 is also available for verification purposes [11].

## V. Methodology

The Bayes rule is at the root of the methodology employed to calibrate the aerodynamic low-order model. In the following, we denote  $\mathbf{q}$  the vector of the parameters subject to calibration i.e., the 8 coefficients of the Bézier curves detailed in Sec. III and  $\mathbf{o}$  the vector of the available observations. According to these definitions, Bayes' rule can be expressed

$$\mathcal{P}(\mathbf{q} | \mathbf{o}) = \frac{\mathcal{P}(\mathbf{o} | \mathbf{q}) \mathcal{P}(\mathbf{q})}{\mathcal{P}(\mathbf{o})}. \quad (8)$$

In (8), we denote  $\mathcal{P}(\mathbf{q} | \mathbf{o})$  the so-called posterior probability distribution of the parameters vector  $\mathbf{q}$ ,  $\mathcal{P}(\mathbf{q})$  the so-called prior probability distribution of the parameters vector  $\mathbf{q}$ ,  $\mathcal{P}(\mathbf{o})$  the evidence of the data, and  $\mathcal{P}(\mathbf{o} | \mathbf{q})$  the so-called likelihood of the data.

The Bayesian approach requires the encoding of all the prior knowledge concerning  $\mathbf{q}$  into  $\mathcal{P}(\mathbf{q})$ . This task translates into shaping the prior probability distribution around specific values, according to what experience or previous knowledge suggests. Having defined the prior, the Bayes rule provides a means to update the distribution of  $\mathbf{q}$  after including new data available from the observation of an event ( $\mathbf{o}$ ). We stress here that the explicit computation of the evidence is not necessary for our purposes. Indeed, we are interested in obtaining calibrated values i.e., the most probable values after including observations, whereas the evidence acts just as a scaling constant normalizing the integral of the posterior distribution to 1. Obtaining the posterior distribution requires also the likelihood  $\mathcal{P}(\mathbf{o} | \mathbf{q})$  to be specified. The likelihood is the probability of observing a particular datum  $\mathbf{o}$ , according to the model and given a specific choice of the  $\mathbf{q}$  vector.

In our calibration framework, we assume a limited amount of prior knowledge concerning the model values providing the best fit to data, being an upper and a lower bound the only information imposed by plausibility constraints. For this reason, our priors are uniform distributions within a finite range  $Q$ , disregarding any correlation or dependence between the parameters in  $\mathbf{q}$ . Namely, we assume that  $\mathbf{q} \in Q$ , being  $Q$  a list of ranges for the acceptable values for each of the free parameters in the low-order model. According to this assumption, which is convenient from the computational

point of view, the resulting posterior is practically proportional to the likelihood for  $\mathbf{q} \in Q$  and vanishing for  $\mathbf{q} \notin Q$ <sup>‡</sup>.

$$\mathcal{P}(\mathbf{q} | \mathbf{o}) \propto \begin{cases} \mathcal{P}(\mathbf{o} | \mathbf{q}) & \mathbf{q} \in Q, \\ 0 & \mathbf{q} \notin Q. \end{cases} \quad (9)$$

Concerning the likelihood, we implement an approach coherent with the principle of maximum entropy [14], assuming the likelihood to be a product of Gaussian distributions, each centered on a specific single observation  $o_i$  with a known variance of  $\sigma_i^2$ .

$$\mathcal{P}(\mathbf{o} | \mathbf{q}) = \prod_i \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left[-\frac{|o_i - \bar{\delta}(\mathbf{q})|^2}{2\sigma_i^2}\right]. \quad (10)$$

In practice, values of  $\mathbf{q}$  yielding large discrepancies (between the  $\mathbf{o}$  and their prediction  $\bar{\delta}(\mathbf{q})$  obtained by evaluating the CFD model) are penalized. The primary hypothesis underlying Eq. (10) consists in assuming the independence of the error on each of the components of  $\mathbf{o}$ . The independent error model is justified in the case of experimental measurement noise such as the one we are interested in.

The Bayes' rule, therefore, reduces to a combination of the information available prior to the experiment, encoded in  $\mathcal{P}(\mathbf{q})$ , with the newly collected data, through the likelihood  $\mathcal{P}(\mathbf{o} | \mathbf{q})$ . Unfortunately, the computation of this latter requires the evaluation of the full CFD model which, of course, is implemented into complex software. As a consequence, there is no analytic approach for obtaining the posterior, which requires instead the application of brute force approaches or efficient sampling methods. A brute force approach, or grid search method, consists in evaluating the posterior numerically on a regularly spaced grid spanning the full prior probability space with the desired resolution. As a drawback, the cost of this grid search explodes with the dimensions of the prior probability space. Markov Chain Monte Carlo (MCMC) sampling methods provide instead a valid alternative when dealing with applications entailing complex multi-variate probability distributions. Here, we rely on the popular Metropolis-Hastings (MH) algorithm [15] which consists in drawing a large sequence of samples, typically tens to hundreds of thousands, from the posterior distribution. Basically, the method implements an efficient random walk guided by an accept/reject criteria. Once a sufficient number of samples has been collected, possibly less than what is required from a brute approach method, it is then possible to reconstruct the whole distribution.

In any case, the computation of the likelihood of a sample requires the evaluation of the full CFD model for obtaining  $\bar{\delta}(\mathbf{q})$ . Unfortunately, even if the evaluation of the low-order computational model is still considered cheap from the perspective of CFD (the evaluation for the considered problem requires simulations from the order of seconds to a few minutes), the calibration task is still too demanding since it requires hundreds of thousands of evaluations. Therefore, a surrogate of the CFD model is necessary and must be devised.

### A. The surrogate model

To reduce the cost of the calibration process, we build a surrogate of the DUST computational model based on a Polynomial Chaos (PC) expansion [16, 17]. Namely, we built a metamodel to mimic the mapping of the model parameters ( $\mathbf{q}$ ) to the QoI ( $\bar{\delta} = \bar{\delta}(\mathbf{q})$ , a generic output of the full computational model corresponding to the observations available later for the calibration).

Assuming  $\mathbf{q} \in \Omega^d$ , with  $q_i \sim \mathcal{U}(q_i^{\min}, q_i^{\max}) \forall i = 1, \dots, d$ , the expansion corresponds to the orthogonal projection, into the linear span, of a finite set of orthonormal polynomials  $\Psi_k(\mathbf{q})$

$$\bar{\delta}(\mathbf{q}) \doteq \sum_{k=0}^{\infty} c_k \Psi_k(\mathbf{q}), \quad (11)$$

where  $c_k$  are unknown coefficients and where the index of summation  $k$  indicates the polynomial order. In particular,  $k$  ranges from 0 to  $\infty$ . In practice, an upper arbitrary value  $P$  must be set to truncate the expansion. In this paper, we rely on the so-called *hyperbolic* truncation designed by Blatman and Sudret [18].

The polynomial basis  $\Psi(\mathbf{q})$  is chosen such that the following orthogonality relation, corresponding to the inner product in the  $L_2$  Hilbert space, is satisfied

$$\langle \Psi_i(\mathbf{q}), \Psi_j(\mathbf{q}) \rangle = \int \Psi_i(\mathbf{q}) \Psi_j(\mathbf{q}) p(\mathbf{q}) d\mathbf{q} = \langle \Psi_i(\mathbf{q})^2 \rangle \delta_{ij}. \quad (12)$$

<sup>‡</sup>we refer to the fact that, in the calibration process, we consider the logarithm of the posterior. In the logarithm form, the uniform prior assumption results in an additive constant term which can be neglected.

Recalling the assumption of that  $q_i \sim \mathcal{U}(q_i^{\min}, q_i^{\max}) \forall i = 1, \dots, d$ , and according to the Askey scheme, see Ref. [19], the polynomial basis can be generated by *Legendre* orthogonal polynomials. The projection coefficients  $\mathbf{c}$  can be computed according to different methodologies. In this work, the expansion coefficients are computed using a non-intrusive regressive approach based on a least-square regression. Namely, once the truncation scheme has been selected, we solve the discrete minimization problem

$$\mathbf{c} = \arg \min_{\mathbf{c} \in \mathbb{R}} \left\{ \frac{1}{N} \sum_{n=1}^N (\tilde{\delta}(\mathbf{q}_n) - \sum_{k=0}^P c_k \Psi_k(\mathbf{q}_n))^2 \right\}, \quad (13)$$

where the  $\mathbf{q}_n$  with  $n = 1 \dots N$  is a set of training points generated using a Latin Hypercube Sampling (LHS) technique. Once a sufficient amount of realizations is available, the expansion coefficients of the linear regression can be reasonably approximated.

## VI. Results

In this section, we present results concerning the calibration of the low-order model. First, we present the accuracy assessment of the surrogates employed to hasten the calibration process and the verification of the calibration framework. After, we present the results of the calibration procedure and the improvements achieved in terms of accuracy gain of predictions from the DUST solver.

### A. Code verification

We first produce an analysis of the accuracy of the surrogates models employed to hasten the calibration procedure. After we assess the performances of the calibration algorithm w.r.t. a dummy test case, ensuring that the overall framework is capable of retrieving a preset target solution  $\hat{\mathbf{q}}$ .

#### 1. Surrogate assessment

In our framework, we deal with several configurations of the wing in the application case. Namely, we consider a wide range of angles of attack and of deflection of the control surface. Since these configurations are treated as discrete points within a well defined range of values of practical interest, the test campaign results in data associated to all the possible combinations of the  $\alpha - \delta$  points. Therefore, a dedicated surrogate is built for each combination of the  $\alpha - \delta$  angles, considering 8 input parameters corresponding to the unknown parameters to be inferred. Figure 2 reports the analysis of the performance of the trained PC w.r.t. both the training data set (TRN), including 600 points, and the verification data set (VER), including 100 points<sup>§</sup>. Namely, the two plots report the root mean squared deviation (RMSD) and the maximum absolute discrepancy in between the surrogate prediction and the reference data point.


$$\text{RMSD} = \sqrt{\frac{1}{N} \sum_{n=1}^N |\tilde{\delta}(\mathbf{q}_n) - \sum_{k=0}^P c_k \Psi_k(\mathbf{q}_n)|^2}, \quad \text{and} \quad \text{MAX} = \max \left( \left| \tilde{\delta}(\mathbf{q}_n) - \sum_{k=0}^P c_k \Psi_k(\mathbf{q}_n) \right| \right), \quad \text{with } n = 1, \dots, N.$$

Each plot features vertical shaded bands which help identifying the angle of attack corresponding to the surrogate model. Moreover, within each band, the surrogates are ordered from the lowest (-35 degrees) to the largest (+35 degrees) deflection of the control surface.

The accuracy assessment reveals a not very high quality of surrogates. Figure 2a shows that the trained surrogates fail in reproducing all the training values. On average, the RMSD is notable but not too dramatic, being always smaller than 0.1. On the other hand, the MAX discrepancy may achieve significant values of 0.2. A similar behavior is observed considering the validation set, in Fig. 2b, from which we also note the poor quality of some surrogates corresponding to particular wing configuration, or in other words combinations of the  $\alpha - \delta$  values.

The reason for such not very satisfactory quality is to be found in the poor stability of the implemented sweep correction with respect to very large variations considered for the parameters included in the  $\mathbf{q}$  vector. This is surely a key aspect, and future improvements should aim at solving this issue. Despite this fact, we will show later that the calibration procedure succeeds at inferring  $\mathbf{q}$ , leading to improving the predictions from the low order model.

<sup>§</sup>Actually, 600 TRN points and 100 VER points for each  $\alpha - \delta$  configuration, for a total of 117'600 TRN points and 19'600 VER points.



**Fig. 2** Assessment of the PC surrogates, reporting the root mean squared deviation (RMSD) and the maximum of the squared error (MAX). (a) Assessment w.r.t. the training set TRN. (b) Assessment w.r.t. the training set VER.

**Table 2** Target values ( $\hat{\mathbf{q}}$ ), and values inferred using the MLE ( $\mathbf{q}_{MLE}$ ) and the MCMC ( $\mathbf{q}_{MCMC}$ ) obtained for the verification test case. The table reports also the prior search ranges.

	TIP				ROOT			
	$q_1$ [0.0, 2.0]	$q_2$ [0.0, 2.0]	$q_3$ [0.0, 2.0]	$q_4$ [0.0, 2.0]	$q_5$ [-2.0, 0.0]	$q_6$ [0.0, 2.0]	$q_7$ [0.0, 2.0]	$q_8$ [0.0, 2.0]
$\hat{\mathbf{q}}$	0.400	1.000	1.000	1.600	-1.000	0.600	0.000	0.200
$\mathbf{q}_{MLE}$	0.399	1.000	1.000	1.599	-1.001	0.600	0.001	0.199
$\mathbf{q}_{MCMC}$	0.312	1.024	1.076	1.560	-1.012	0.678	0.006	0.190

## 2. MCMC algorithm verification

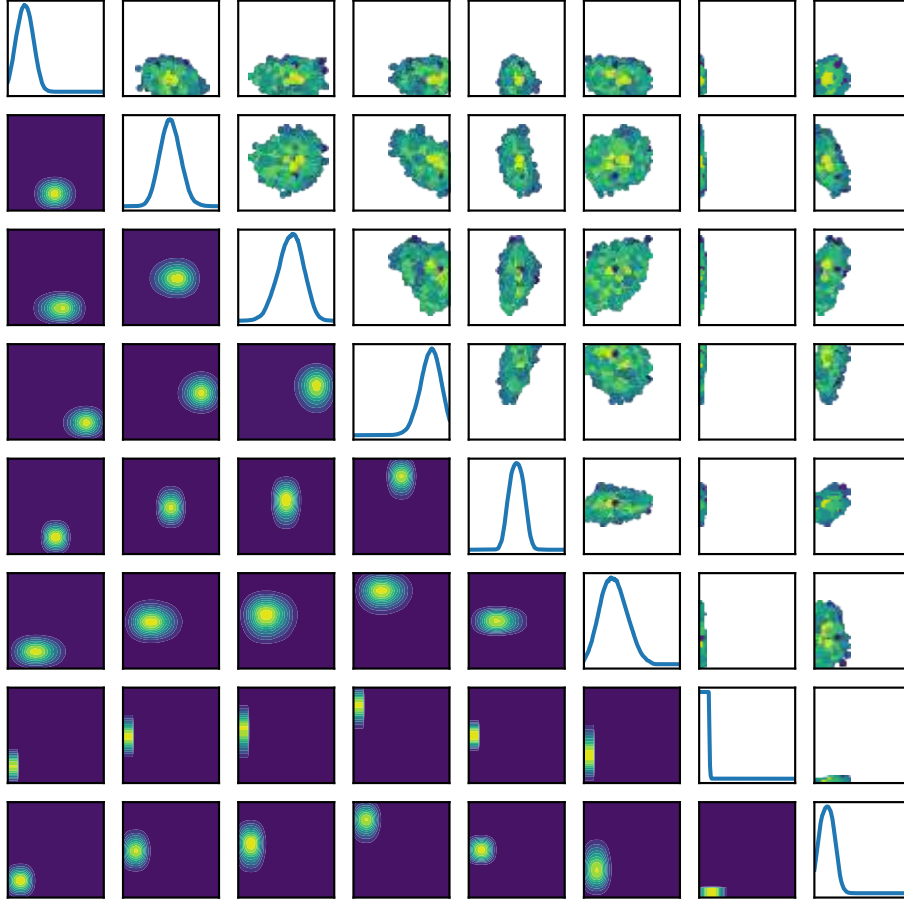
This section verifies the implementation of the MCMC algorithm by assessing its capacity to correctly infer arbitrary target coefficients  $\hat{\mathbf{q}}$ , see Tab. 2. In this context, the verification is carried out by exploiting a dummy test case entailing data produced by evaluating the surrogate models at  $\hat{\mathbf{q}}$ . The target vector entries are reported in Tab. 2 together with the results from the calibration algorithm. In particular,  $\mathbf{q}_{MLE}$  (Maximum Likelihood Estimation) reports the values obtained by the mere application of an optimization algorithm to the likelihood function, whereas  $\mathbf{q}_{MCMC}$  reports the maximum a posteriori values resulting from the deployment of the MCMC chain. As we are dealing with a dummy test case, we selected an arbitrary  $\sigma_i = 0.02$ , equal for all data points. Clearly, both the MLE and the MCMC approaches are reasonably capable of returning the targeted values, supporting the verification of the calibration framework.

The specific set up of the MCMC algorithm considers a burn-in phase of  $10^4$  samples and a MCMC chain of  $150^4$  samples. The chain is endowed with a very good mixing and a very fast decay of all the auto-correlation functions for all the eight parameters (not reported). Although the MCMC acceptance rate is quite low (about 4%), the indicators assessing the convergence of the posteriors reach stationary values well within a few thousands of samples, a number considerably lower than the overall MCMC chain length. These indicators are: the means of the marginal posteriors, their standard deviation, and the  $L_1$  norm of the kernel density estimates obtained considering different chain portions.

Figure 3 reports the resulting posterior marginal PDFs on the diagonal plots, showing that the calibration framework leads to the identification of regions of very high likelihood for each of the eight unknown parameters, with the bulk of the probability mass distributed in the proximity of the target values. Moreover, Figure 3 reports also the joint marginal



distributions, on the lower triangular side, and the sampled points projected on the 2D space with color proportional to the likelihood, on the upper triangular side.



**Fig. 3** Verification test case using dummy data. Picture reports the posterior marginal PDFs on the diagonal plots. On the extra diagonal entries: on the lower side the joint marginal distributions are reported, whereas on the upper side are the sampled points projected on the 2D space with color proportional to the likelihood.

Concerning the MLE approach, the optimization method is the Limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS-B) algorithm, a popular choice for parameter estimation problems in machine learning [20, 21]. Note that we implemented an optimization procedure entailing 10 independent search processes, each based on a random initialization point. Initializing the optimization from different random points reduces the risk of obtaining a solution associated to a region of local minimum. Of 10 runs, we obtained the targeted solution (which is anyway the best solution) for 5 different and random initialization points.

### B. Model calibration

In this Section, we present the results of the calibration procedure carried out considering three different data sets. Hereinafter, we establish the following labeling for the identification of the considered data sets: RANS is the data set including CFD data only, EXP is the data set including experimental data only, whereas MIXED is the merged set of both CFD and experimental data. In any case, 20% of the data points are retained for validation purposes namely, to

**Table 3** Calibrated values ( $\mathbf{q}$ ) obtained w.r.t. the different data sets. The table reports also the prior search ranges.

	TIP				ROOT			
	$q_1$ [0.0, 2.0]	$q_2$ [0.0, 2.0]	$q_3$ [0.0, 2.0]	$q_4$ [0.0, 2.0]	$q_5$ [-2.0, 0.0]	$q_6$ [0.0, 2.0]	$q_7$ [0.0, 2.0]	$q_8$ [0.0, 2.0]
RANS	1.272	0.093	0.172	0.029	-1.988	1.091	0.278	0.018
EXP	1.549	0.293	1.856	0.129	-0.106	0.176	0.035	1.644
MIXED	0.182	0.074	0.011	0.399	-1.990	0.863	0.098	0.091

assess the accuracy of the low-order model after calibration.

For all the considered cases, the calibration process entails a burn-in phase of 10'000 samples and a total length of the MCMC chain of 150'000 samples. We also consider a  $\sigma_i = 0.2$ , equal for all points, irrespective of the data source. Note that, in general, this latter assumption may be changed if one desires to assign a different weight to data produced from a specific source.

Table 3 reports the values of the model parameter after calibration w.r.t. the different data sets. A first comparison reveals that results change significantly, depending on the considered data source. In the following, we provide a deeper analysis of the results obtained for each data set.

#### 1. Calibration using RANS data

Data are provided from [11]. Three-dimensional CFD simulations of the cantilever wing in a free stream were carried out employing the RANS model from the SU2 open-source suite [13] coupled to the Spalart-Allmaras (SA) [22] closure. SU2 is a computational tool for solving Partial Differential Equations (PDEs) on unstructured grids and it implements a Finite Volume Method using a standard edge-based data structure on a vertex-based dual grid.

Steady-state simulations are performed using an implicit Euler time-marching scheme. The convergence to a steady-state solution is assessed by evaluating the steadiness of the predicted lift ( $C_L$ ) and drag ( $C_D$ ) coefficients. The convective fluxes are discretized using an upwind Roe scheme with a Monotonic Upstream-centered Scheme for Conservation Laws (MUSCL) for achieving a second-order accuracy. The Venkatakrisnan flux limiter is employed to damp local solution oscillation that may possibly arise in domain regions characterized by strong gradients.


Because of the large number of different geometries subject to analysis, an automatic meshing procedure has been devised. Starting from the specification of the parameters defining the specific configuration of the wing, a parametric CAD model is built using the software CATIA V5. After, a set of dedicated scripts is used to transfer the CAD model to the meshing software Pointwise [23]. For the surface mesh, highly-anisotropic triangular cells are used to increase resolution in high-curvature regions, such as the leading edge and the tail-rudder junction surface. For the volume mesh, a hex-dominant, octree-based mesh is built to limit numerical dissipation. Both the surface and the volume meshes are augmented through mesh refinement in regions where complex flow patterns are expected e.g., the portion of the surface in close proximity to the tip region. An  $\alpha$ -adaptive wake refinement is also performed in the direction of the freestream flow, to better capture the wake developing past the wing.

The original data set consists of 420 simulation points spanning the experimental conditions reported in Sec. IV. In this Section, we consider only the data corresponding to a null dihedral angle ( $\Gamma = 0$ ) deg, thus reducing the number of available data points. In total, we consider a subset including 206 points. Namely, 164 points are employed for the calibration procedure, whereas 42 points are employed for the a posteriori accuracy assessment of the calibrated low-order model.

Table 3 reports the values inferred using our calibration framework ( $\mathbf{q}$ ), and the search intervals for each parameter. The chain shows a good mixing and a fast decay, within about 1000 samples, of the auto-correlation function associated with each parameter (not reported for brevity), therefore confirming the suitable setting of the sampling algorithm. The acceptance rate of the chain after discarding the burn-in portion is about 2.5%, therefore a quite low value. Anyways, the resulting posteriors are converged w.r.t. the number of chain samples, achieving a stationary distribution after about 80'000 samples. Instead, the mean and the standard deviation of marginal posteriors are stationary after about 40'000 samples.

Figure 4 reports the posterior marginal PDFs on the diagonal plots, the joint marginal distributions on the lower triangular side, and the sampled points projected on the 2D space with color proportional to the likelihood on the upper

triangular side. All posteriors show a concentration of probability mass in specific regions of the admissible range. Moreover, some correlation may be observed among of parameter pairs e.g.,  $q_7 - q_8$ . Practically, only the posterior



**Fig. 4 Calibration test case using RANS data. Picture reports the posterior marginal PDFs on the diagonal plots. On the extra diagonal entries: on the lower side the joint marginal distributions are reported, whereas on the upper side are the sampled points projected on the 2D space with color proportional to the likelihood.**

associated to the  $q_6$  parameter presents a peak well included between the prior range. The remaining parameters present the bulk of the posterior probability distribution backed against either the upper or the lower prior bound, indicating that the selected prior bounds are likely too narrow. Unfortunately, such a narrow priors had to be selected in order to ensure the stability of the numerical solver implemented in DUST, and therefore to guarantee the reliability of low-order simulations and, additionally, to achieve surrogates of acceptable quality. Nonetheless, the calibration process succeeds, ultimately leading to an improvement of the low-order model predictions.

In Fig. 5a we report the comparison of predictions from the standard DUST solver (STD), prediction from the DUST solver implementing the correction with calibrated values ( $COR_{MCMC}$ ), and the validation data set (Data). The plot gathers the validation points corresponding to the same value of the angle of attack  $\alpha$ , but with a different value of the control surface deflection  $\delta$ . As shown in the figure, there exists a gap between the predictions from the standard solver and the RANS data. This gap is due to the differences inherent the low-order model and the fully 3D RANS model. Thanks to the implementation of the sweep angle correction, and its subsequent calibration, the gap is reduced

and this is evident in particular in the region of large  $\alpha$ , namely  $\alpha > 8$  degrees.

This fact is highlighted in Fig. 5b where we plot the gain/loss attained in terms of discrepancy w.r.t. the validation data points. Namely, we compute the quantity  $|C_L^{STD} - C_L^{Data}| - |C_L^{COR} - C_L^{Data}|$  which is positive (green) if the calibrated correction provides a reduction of the discrepancy, and therefore the corrected model performs better than the standard one, negative otherwise (red). As clear from the plot, the calibrated corrected model leads to a significant improvement of the prediction accuracy throughout the whole span of the  $\alpha - \delta$  domain, with just a couple of points revealing a degradation.

## 2. Calibration using EXP data

The experimental data set was generated through two experimental test campaigns carried out at the "Galleria del Vento del Politecnico di Milano" (GVPM), a closed-circuit wind tunnel, arranged in a vertical layout with two test sections located on the opposite sides of the loop. The considered test campaigns were performed in the Low-Turbulence Test Section, capable of reaching a maximum wind speed of 55 m/s and a turbulence level less than 0.1%, see Ref. 24. The specific set up of the experiment and its details can be found in Ref. 10.

From the full experimental database, we select the measurements corresponding to a null dihedral angle for a total of 51 points. Again, we retain 20% of points (11) for validation purposes and use the remaining 80% (40) to carry out the calibration.

Table 3 reports the values inferred using our calibration framework ( $\mathbf{q}$ ), and the search intervals for each parameter. The setting of the calibration process is the same employed in RANS case. Though the chain shows a fair mixing, it fails in clearly identifying a region of very high likelihood for the eight unknown parameters. The prior domain is thoroughly and continuously explored, resulting in quite poorly informative posteriors. The auto-correlation functions associated with the parameters show a very fast decay (in the order of a few hundreds samples, with an acceptance rate of about 6%). Anyways, the indicators considered for assessing the convergence of the posteriors namely, the posterior means, standard deviations, and the  $L_1$  norm of the kernel density estimates obtained considering different chain portions, reach stationary values well within a few thousands of samples, a number considerably lower than the overall MCMC chain length.

Figure 6 plots the posterior marginal PDFs on the diagonal. Clearly, posteriors are in general poorly informative, except for  $q_5$ ,  $q_7$ , and perhaps  $q_8$ . In particular, the posteriors associated to  $q_1$ ,  $q_2$  and  $q_4$  are very flat, indicating an uniform probability distribution of such parameters. At the same time, this facts are reflected onto the joint marginal distributions reported on the lower triangular side, or by looking at the sampled points projected on the 2D space, with color proportional to the likelihood, reported on the upper triangular side. Practically, the amount of information extracted by the reference data set is not sufficient for clearly identifying a particular value of many of the unknown parameters, except indeed  $q_5$ ,  $q_7$ . Nonetheless, we select the maximum a posteriori point from the chain, and carry out the validation which is reported in Fig. 7a. Again, we compare predictions from the standard DUST solver (STD) with prediction from the corrected DUST solver exploiting this new calibrated values ( $COR_{MCMC}$ ), and the validation data set (Data).

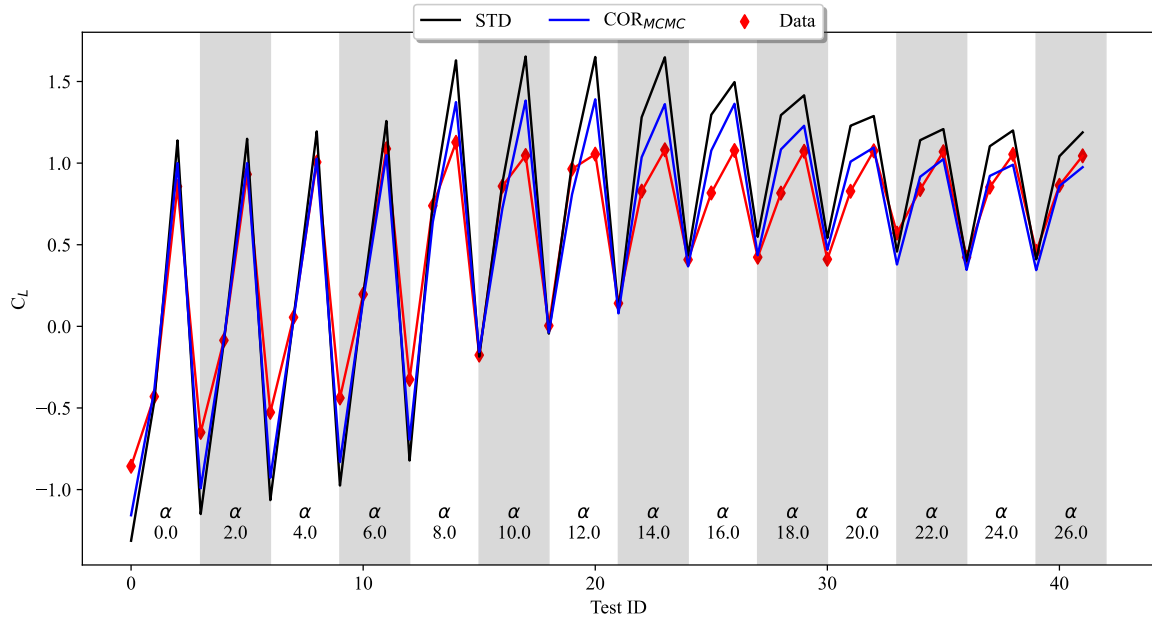
Despite the scarcity of information brought by the EXP data set, the prediction gap is varied thanks to the calibrated sweep angle correction. Improvements are surely limited to well specific regions and, differently than the RANS case, are now counterbalanced by a significant loss of accuracy for some test configurations. This fact can be better appreciated in Fig. 7b where we plot the gain/loss attained in terms of discrepancy w.r.t. the validation data points.

## 3. Calibration using MIXED data


The MIXED reference data set is composed by the merging of both the RANS and the EXP data set. Namely, a total of 257 points is now available, with 205 reserved for the calibration and 52 for validation of the results.

The inferred values are reported in Tab. 3 and show some similarity w.r.t. the values obtained using the RANS data set. This was expected since the MIXED data set contains a large majority of CFD data. In particular,  $q_2$ ,  $q_5$  and  $q_8$  are quite similar to the values obtained using the RANS data set only. On the other hand, the inferred  $q_1$  and  $q_4$  are quite different, whereas more contained discrepancies remains for the remaining parameters, indicating that the contribution of the information encoded in the EXP data is not negligible.

The calibration results are reported in Tab. 3. For this test case, the autocorrelation decays quickly along the chain and a good mixing of the eight parameters is detected, with a poor acceptance rate of about 2.3%. Anyways, the indicators considered for assessing the convergence of the posteriors reach stationary values well within the overall length of the MCMC chain.




(a) RANS data set. Predictions comparison.



(b) Gain/loss plot.

**Fig. 5 Calibration test case using RANS data. (a) Comparison of predictions from the standard DUST solver (STD), prediction from the DUST solver implementing the correction with calibrated values ( $COR_{MCMC}$ ), and the validation data set (Data). (b) Accuracy gain/loss plot highlighting the improvements achieved using the calibrated correction, in terms of  $\Delta C_L$  w.r.t. the validation data set.**

Figure 8 reports the usual table plot concerning the calibration procedure. Similarly to the RANS case, there are regions of high likelihood which are clearly identified as the chain explores the prior probability space. For almost all



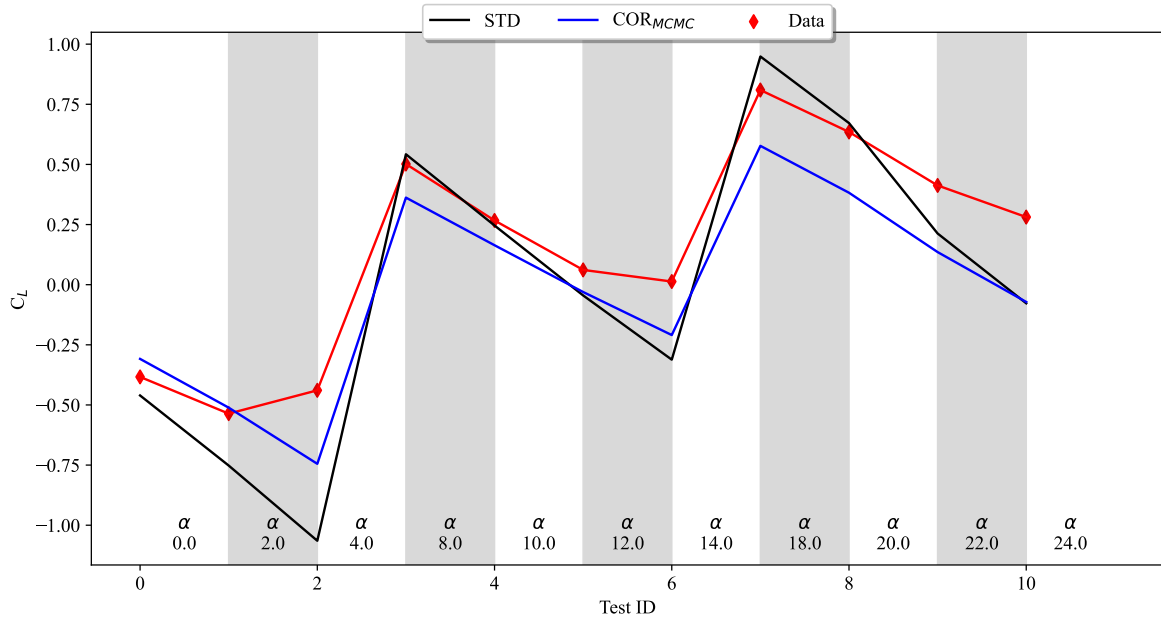
**Fig. 6 Calibration test case using EXP data. Picture reports the posterior marginal PDFs on the diagonal plots. On the extra diagonal entries: on the lower side the joint marginal distributions are reported, whereas on the upper side are the sampled points projected on the 2D space with color proportional to the likelihood.**

parameters, the bulk of the posterior probability distribution is backed against one of the prior bounds, stressing out again that the selected prior bounds are too narrow and that better results may be achieved if one considers a wider prior range. As mentioned, this is currently not feasible since it entails a loss of numerical stability concerning the computation of the aerodynamic solution by means of the DUST solver.

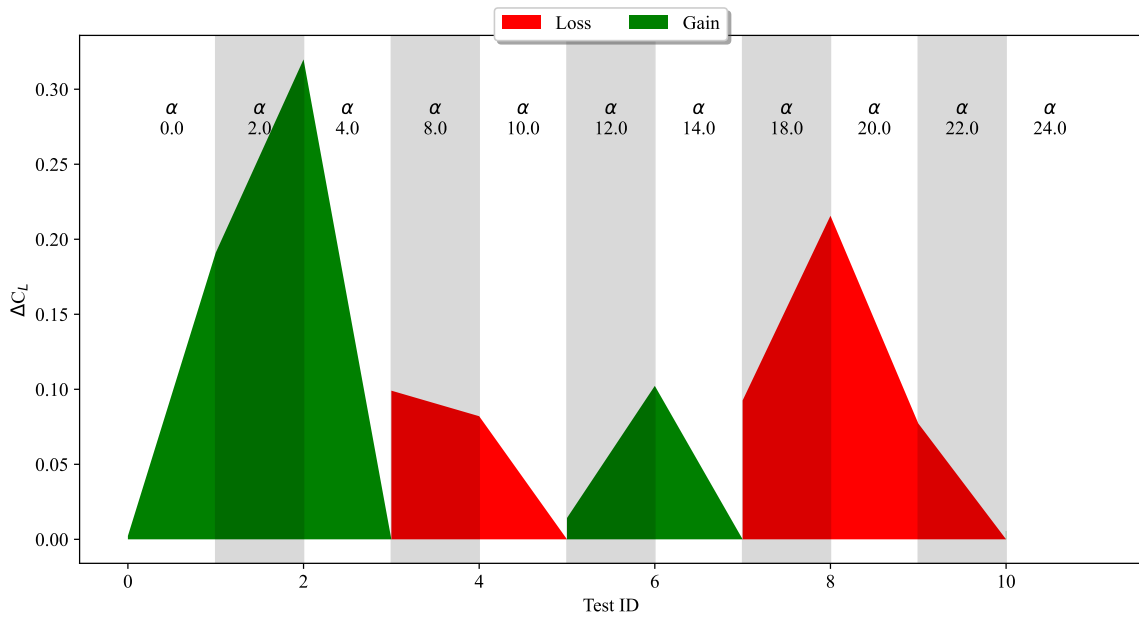
Despite the narrow priors, we select the maximum a posteriori point and use the corresponding parameter values to carry out the model accuracy assessment, in Fig. 9a. Not surprisingly, and because of the composition of the MIXED data set, conclusions are similar to the ones drawn from the RANS case. The noticeable difference, which is clear from Fig. 9b is that the attained gain is larger (about 0.4 in place of about 0.3 achieved using RANS data only). At the same time, we also note the occurrence of additional regions in which there is actually a loss of accuracy, although limited.

## VII. Conclusions

A Bayesian calibration framework was devised to improve numerical predictions from an aerodynamic low-order model with regards to a specific application case of interest for the MONNALISA project, which develops in the frame of the CS2 Joint Undertaking. Namely, a cantilever trapezoidal wing in a free stream flow, which represents the horizontal




(a) EXP data set. Predictions comparison.



(b) Gain/loss plot.

**Fig. 7 Calibration test case using EXP data. (a) Comparison of predictions from the standard DUST solver (STD), prediction from the DUST solver implementing the correction with calibrated values ( $COR_{MCMC}$ ), and the validation data set (Data). (b) Accuracy gain/loss plot highlighting the improvements achieved using the calibrated correction, in terms of  $\Delta C_L$  w.r.t. the validation data set.**

component of the rear end of a commercial airplane. Since the goal of the MONNALISA project is to devise tools and knowledge to support the development of an unconventional “Advanced Rear End” component for the forthcoming



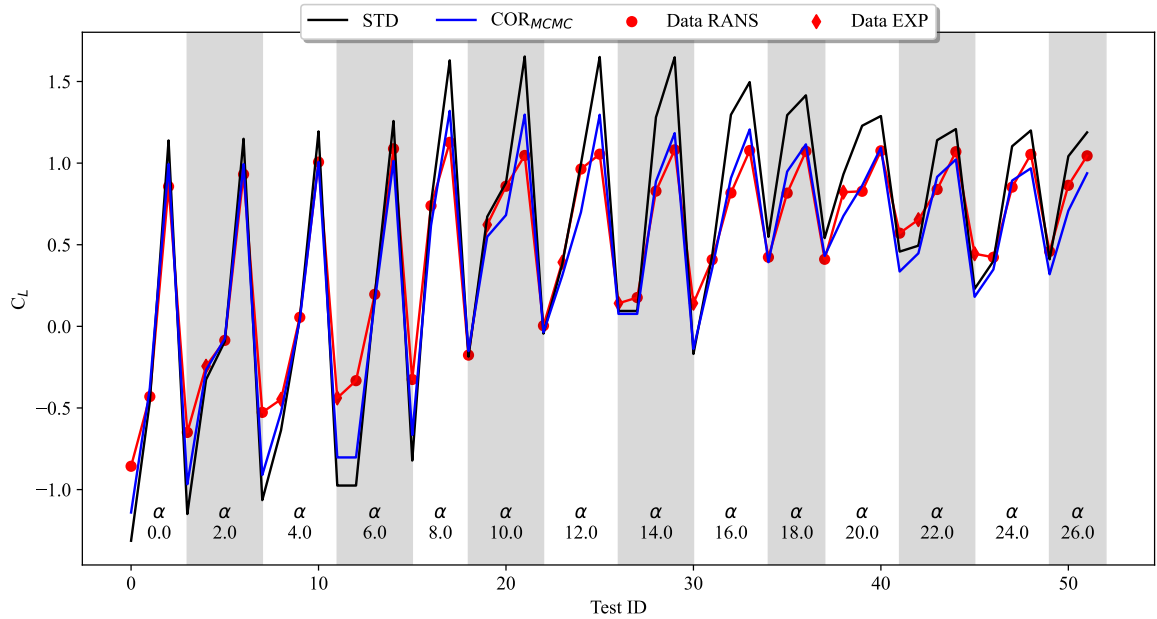
**Fig. 8 Calibration test case using MIXED data. Picture reports the posterior marginal PDFs on the diagonal plots. On the extra diagonal entries: on the lower side the joint marginal distributions are reported, whereas on the upper side are the sampled points projected on the 2D space with color proportional to the likelihood.**

generation of ultra-efficient aircraft, the application case is investigated considering a multitude of operating conditions. Namely, high incidence angles resulting into a stalled flow, or different dihedral and sweep angles and taper ratios.

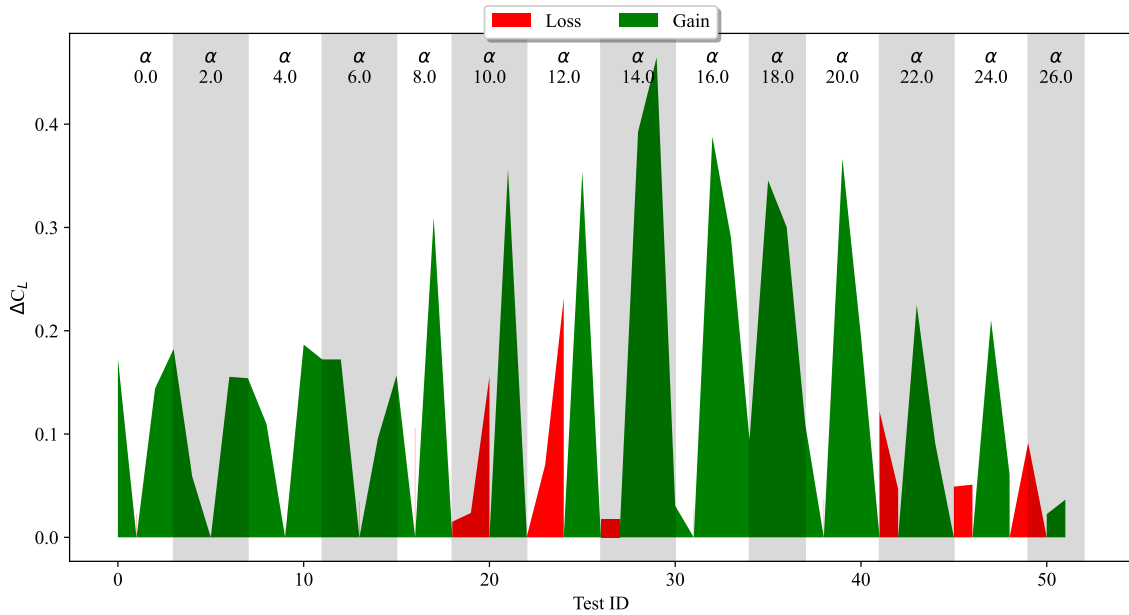
In this work, it is shown that the formulation of the calibration procedure, its implementation into a computer code, and its setting, are suitable for achieving a significant improvement of predictions from the aerodynamic low-order model, including also predictions concerning the post-stalled regime. After, it is reported that substantial improvements were obtained with respect to the data available from RANS simulations, from experiments, and also for a mixed data set. The improvements are shown for a number of different operating conditions entailing a large variation of the angle of attack and of the deflection of the control surface, but limited to a null dihedral angle, a back-sweep of 30 degrees, and fixed taper ratio, and fixed Mach and Reynolds number.

Nonetheless, there remain regions of the operative conditions domain for which further improvements may be possible and necessary. Future work may be devoted to enriching the calibration data set with the inclusion of data obtained with a diverse operating conditions e.g., a non null dihedral. Moreover, a possible future development may consist in substituting the cubic Bézier curves (appearing in the sweep angle correction to the low order model) with more flexible functionals.





(a) MIXED data set. Predictions comparison.



(b) Gain/loss plot.

**Fig. 9** Calibration test case using MIXED data. (a) Comparison of predictions from the standard DUST solver (STD), prediction from the DUST solver implementing the correction with calibrated values ( $COR_{MCMC}$ ), and the validation data set (Data). (b) Accuracy gain/loss plot highlighting the improvements achieved using the calibrated correction, in terms of  $\Delta C_L$  w.r.t. the validation data set.

### Acknowledgments

This work is part of the MONNALISA project (<https://www.monnalisa-project.eu>) which has received funding from the Clean Sky 2 Joint Undertaking (JU) under grant agreement No 101008257. The JU receives support from the

European Union’s Horizon 2020 research and innovation programme and the Clean Sky 2 JU members other than the Union. All the material reported in the present paper reflects the view of the authors only. The JU is not responsible for any use that may be made of the information it contains.

G. Gori would like to acknowledge that this work was partially funded by the European Union (Project 101059320 - UN-BIASED). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Executive Agency. Neither the European Union nor the European Research Executive Agency can be held responsible for them.

## References

- [1] Brouckaert, J.-F., Mirville, F., Phuah, K., and Taferner, P., “Clean Sky research and demonstration programmes for next-generation aircraft engines,” *The Aeronautical Journal*, Vol. 122, No. 1254, 2018, p. 1163–1175. <https://doi.org/10.1017/aer.2018.37>
- [2] Auteri, F., Gibertini, G., Gori, G., Guardone, A., Rausa, A., Zanotti, A., Congedo, P. M., Menzago, A., and Llamas-Sandin, R. C., “Monnalisa: Modelling Nonlinear Aerodynamics of Lifting Surfaces,” *AIAA AVIATION 2022 Forum*, 2022. <https://doi.org/10.2514/6.2022-4149> URL <https://arc.aiaa.org/doi/abs/10.2514/6.2022-4149>
- [3] Tugnoli, M., Montagnani, D., Syal, M., Droandi, G., and Zanotti, A., “Mid-fidelity approach to aerodynamic simulations of unconventional VTOL aircraft configurations,” *Aerospace Science and Technology*, Vol. 115, 2021, p. 106804. <https://doi.org/https://doi.org/10.1016/j.ast.2021.106804> URL <https://www.sciencedirect.com/science/article/pii/S127096382100314X>
- [4] Dussin, D., Fossati, M., Guardone, A., and Vigeveno, L., “Hybrid grid generation for two-dimensional high-Reynolds flows,” *Computers & Fluids*, Vol. 38, No. 10, 2009, pp. 1863–1875.
- [5] Goitia, H., and Llamas, R., “Nonlinear vortex lattice method for stall prediction,” *MATEC Web of Conferences*, Vol. 304, 2019, p. 02006. <https://doi.org/https://doi.org/10.1051/mateconf/201930402006> URL <https://www.matec-conferences.org/10.1051/mateconf/201930402006>
- [6] Hosangadi, P., Paul, R. C., and Gopalarathnam, A., “Improved Stall Prediction for Swept Wings Using Low-Order Aerodynamics,” *33rd AIAA Applied Aerodynamics Conference*, 2015. <https://doi.org/https://doi.org/10.2514/6.2015-3159> URL <https://arc.aiaa.org/doi/abs/10.2514/6.2015-3159>
- [7] Mariens, J., Elham, A., and van Tooren, M. J. L., “Quasi-Three-Dimensional Aerodynamic Solver for Multidisciplinary Design Optimization of Lifting Surfaces,” *Journal of Aircraft*, Vol. 51, No. 2, 2014, pp. 547–558. <https://doi.org/https://doi.org/10.2514/1.C032261> URL <https://doi.org/https://doi.org/10.2514/1.C032261>
- [8] Petrilli, J. L., Paul, R. C., Gopalarathnam, A., and Frink, N. T., “A CFD Database for Airfoils and Wings at Post-Stall Angles of Attack,” *31st AIAA Applied Aerodynamics Conference*, 2013. <https://doi.org/https://doi.org/10.2514/6.2013-2916> URL <https://arc.aiaa.org/doi/abs/10.2514/6.2013-2916>
- [9] Sederberg, T. W., “Computer Aided Geometric Design,” *Computer Aided Geometric Design Course Notes*, 2012. URL <https://scholarsarchive.byu.edu/facpub/1>
- [10] Zanotti, A., Rausa, A., Grassi, D., Riccobene, L., Gibertini, G., Guardone, A., and Auteri, F., “Infrared Thermography Measurements over a Tail-Plane Model of a Large Passenger Aircraft,” *Journal of Physics: Conference Series*, Vol. 2293, No. 1, 2022, p. 012014. <https://doi.org/https://doi.org/10.1088/1742-6596/2293/1/012014> URL <https://dx.doi.org/10.1088/1742-6596/2293/1/012014>
- [11] Rausa, A., Gori, G., Guardone, A., Zanotti, A., and Auteri, F., “Multi-fidelity assessment of the aerodynamic performances of unconventional aircraft tail configurations,” *AIAA AVIATION 2022 Forum*, 2022. <https://doi.org/https://doi.org/10.2514/6.2022-3902> URL <https://arc.aiaa.org/doi/abs/10.2514/6.2022-3902>
- [12] “ERCOFTAC Community,” <https://zenodo.org/communities/ercoftac/about/> 2023.
- [13] Palacios, F., Alonso, J., Duraisamy, K., Colonno, M., Hicken, J., Aranake, A., Campos, A., Copeland, S., Economon, T., Lonkar, A., et al., “Stanford university unstructured (SU2): an open-source integrated computational environment for multi-physics simulation and design,” *51st AIAA aerospace sciences meeting including the new horizons forum and aerospace exposition*, 2013, p. 287.
- [14] Sivia, D. S., and Skilling, J., *Data Analysis - A Bayesian Tutorial*, 2<sup>nd</sup> ed., Oxford Science Publications, Oxford University Press, 2006.

- [15] Hastings, W. K., "Monte Carlo sampling methods using Markov chains and their applications," *Biometrika*, Vol. 57, No. 1, 1970, pp. 97–109. <https://doi.org/10.1093/biomet/57.1.97> URL <http://biomet.oxfordjournals.org/cgi/content/abstract/57/1/97>
- [16] Le Maître, O., and Knio, O., *Spectral Methods for Uncertainty Quantification*, 1<sup>st</sup> ed., Scientific Computation, Springer Netherlands, 2010. <https://doi.org/10.1007/978-90-481-3520-2>
- [17] Reagan, M. T., Najm, H. N., Debusschere, B. J., Le Maître, O. P., Knio, O. M., and Ghanem, R. G., "Spectral stochastic uncertainty quantification in chemical systems," *Combustion Theory and Modelling*, Vol. 8, No. 3, 2004, pp. 607–632. <https://doi.org/10.1088/1364-7830/8/3/010> URL <https://www.tandfonline.com/doi/abs/10.1088/1364-7830/8/3/010>
- [18] Blatman, G., and Sudret, B., "Adaptive sparse polynomial chaos expansion based on least angle regression," *Journal of Computational Physics*, Vol. 230, No. 6, 2011, pp. 2345–2367. <https://doi.org/https://doi.org/10.1016/j.jcp.2010.12.021> URL <https://www.sciencedirect.com/science/article/pii/S0021999110006856>
- [19] Askey, R. A., and Wilson, J. A., "Some basic hypergeometric orthogonal polynomials that generalize Jacobi polynomials," *Memoirs of the American Mathematical Society*, Vol. 54, 1985, pp. 0–0.
- [20] Malouf, R., "A Comparison of Algorithms for Maximum Entropy Parameter Estimation," *Proceedings of the 6th Conference on Natural Language Learning - Volume 20*, Association for Computational Linguistics, USA, 2002, p. 1–7. <https://doi.org/10.3115/1118853.1118871>
- [21] Byrd, R., Lu, P., Nocedal, J., and Zhu, C., "A Limited Memory Algorithm for Bound Constrained Optimization," *SIAM Journal on Scientific Computing*, Vol. 16, 2003. <https://doi.org/10.1137/0916069>
- [22] Spalart, P., and Allmaras, S., "A one-equation turbulence model for aerodynamic flows," *30th aerospace sciences meeting and exhibit*, 1992, p. 439.
- [23] "Pointwise, Inc." <https://www.pointwise.com/> 2007.
- [24] Gibertini, G., Gasparini, L., and Zasso, A., "Aerodynamic design of a civil-aeronautical low speed large wind tunnel," *Proceedings of the AGARD 79th fluid dynamics panel symposium*, 1996.