

Three-dimensional variational data assimilation of separated flows using time-averaged experimental data

Uttam Cadambi Padmanaban,* Bharathram Ganapathisubramani,† and Sean Symon‡

Department of Aeronautics and Astronautics,

University of Southampton,

United Kingdom

(Dated: January 5, 2026)

Abstract

We present a novel framework for assimilating planar PIV experimental data using a variational approach to enhance the predictions of the Spalart–Allmaras RANS turbulence model. Our method applies three-dimensional constraints to the assimilation of mean velocity data, incorporating a corrective forcing term in the momentum equations. The advantages of this approach are highlighted through a direct comparison with traditional two-dimensional assimilation using the same experimental dataset. We demonstrate its efficacy by assimilating the deep stall flow over a NACA0012 airfoil at a 15° angle of attack and a chord-based Reynolds number of $Re_c \approx 7.5 \times 10^4$. We find that in two-dimensional assimilation, the corrective forcing term compensates not only for physical modeling errors but also for the lack of divergence in the experimental data. This conflation makes it difficult to isolate the effects of measurement inconsistencies from deficiencies in the turbulence model. In contrast, three-dimensional assimilation allows the corrective forcing term to primarily address experimental setup errors while enabling the turbulence model to more accurately capture the flow physics. We establish the benefits of three-dimensional assimilation through its ability to account for the inherent three-dimensionality of the flow, thereby enabling more physically consistent reconstructions of key quantities such as pressure, lift force, and eddy viscosity.

* ucp1n22@soton.ac.uk

† g.bharath@soton.ac.uk

‡ sean.symon@soton.ac.uk

I. INTRODUCTION

Data assimilation (DA) first emerged in meteorology [1], where it was used to improve weather predictions by incorporating observational data into numerical models. DA helps reduce observational noise and uncertainties, expand the field of view of available data, and refine modeling assumptions by calibrating model coefficients and boundary conditions. These advantages have led to its increasing use in fluid mechanics, particularly for studying turbulent flows, as demonstrated by recent studies [2–6]. The complexity of turbulence arises from its wide range of length and time scales and its inherently unsteady behavior. Resolving all these scales using direct numerical simulations (DNS) or large eddy simulations (LES) is computationally intractable for high Reynolds number flows of practical interest [7]. A more computationally efficient alternative is the use of Reynolds-Averaged Navier–Stokes (RANS) turbulence models, which rely on modeling assumptions to predict an averaged flow state. However, their accuracy is questionable in cases involving strong pressure gradients and flow separation, as they are typically calibrated against only a limited set of canonical flow problems. The application of DA to fluid mechanics problems is particularly useful in improving RANS turbulence models that are able to capture the effects of turbulence a lot more accurately.

The widespread adoption of DA in fluid mechanics has also been driven by advances in field imaging techniques, such as planar particle image velocimetry (PIV), which provides access to two-component mean velocity data [8]. PIV is often used as an alternative to DNS or LES, particularly for studying high Reynolds number flows of practical interest. However, like any experimental method, PIV is subject to noise and uncertainty. It also comes with additional limitations, including low resolution, a limited field of view, high setup costs (especially due to the laser), large data storage requirements, and long processing times. An important problem that is always overlooked is the lack of the third (out-of-plane) component of velocity which is not imaged using planar PIV, and is non-negligible for separated flows even at low Reynolds number [9, 10]. Assimilating two-component mean velocity data from planar PIV without enforcing three-dimensional (3D) constraints leads to an inaccurate representation of the flow dynamics, as the planar PIV data do not satisfy the 2D continuity equation since the actual flow is three-dimensional.

Among the different types of DA, we use the variational method applied to steady-state

cases in this study. Variational DA optimizes a control variable of choice by minimizing the discrepancy between a model output (such as RANS) and reference data (either from experiments or high fidelity simulation). Gradient-based optimization methods are used to minimize the objective function that quantitatively describes the deviation of model output from reference data. The time-averaged continuity and momentum equations are typically applied as constraints. Adjoint methods are used as an efficient way of computing gradients [11, 12]. The application of 3D DA has been limited due to a variety of reasons. Reference [2] performed variational DA of a flow past a circular cylinder at $Re = 150$ using two-component, mean velocity fields from DNS as reference data. A forcing term in the time-averaged momentum equations was optimized without the application of a turbulence model. The two-dimensional (2D) continuity equation was enforced as a constraint, which is consistent with the 2D nature of the flow [9]. The use of experimental data using a similar methodology for assimilating a higher Reynolds number ($Re_c = 13,500$) flow past an idealized airfoil was demonstrated by Ref. [13]. The justification for enforcing 2D continuity was offered by comparing the optimized forcing with the forcing obtained from the experiment. A good agreement between the two was used to argue that the field was indeed two-dimensional.

Reference [4] extended the work of Ref. [2] by including the Spalart-Allmaras (SA) turbulence model for assimilating the mean velocity field into the RANS equations. This improved the efficiency of the assimilation for high Reynolds number flows as the baseline solution was already a close approximation of the reference mean velocity field. The authors optimized a control variable in the time-averaged momentum equations (similar to Refs. [2, 13]) of a flow past a backward-facing step (BFS) at Reynolds number $Re = 28275$ using DNS reference data. The study utilized a span-averaged and time-averaged velocity field permitting the enforcement of 2D continuity. Additionally, the reference data were obtained from a strictly divergence-free space of mean velocity fields. This meant that the obtained forcing field was only compensating for the lack of physics that the turbulence model was unable to capture and not any deficiency in the quality of reference data. Recently, Ref. [6] used a corrective source term in the turbulence transport equation of the SA model for mean velocity assimilation of a flow past a NACA0012 airfoil. The near-stall flow was at an angle of attack $\alpha = 10^\circ$ and in the Reynolds number range $4.3 \times 10^4 \leq Re \leq 6.4 \times 10^4$. The reference mean velocity field was obtained from time-averaged PIV. DA was able to fill in the gaps that resulted from shadows and reflections during the data acquisition stage of

PIV. The effects of three-dimensionality on reconstructed flow quantities such as Reynolds stresses, nonetheless, were not investigated.

References [2, 13] discussed so far made use of the continuous adjoint approach, which involves linearizing the PDE followed by its discretization. This way, the primal solver used to solve the Navier-Stokes (NS) equations can be re-used for solving the adjoint equations. The discrete adjoint approach, as demonstrated by Ref. [12] and used in Refs. [4, 6], constructs the adjoint equations following its discretization, driving the accuracy of gradient calculation to machine precision. References [14, 15] used a discrete adjoint approach with reference data obtained from LES, which is free from noise, and, more importantly, allows the application of 2D constraints. Reference [16] compared various control variables for 2D mean velocity assimilation, reporting good agreement with reference data and satisfactory Reynolds stress reconstruction.

The need for high-accuracy gradients has led to the emergence of tools that have a discrete adjoint solver implemented in them. DAFoam [17, 18] is one such tool that incorporates a discrete adjoint solver along with primal solvers from OpenFOAM [19]. Since the code is open source, easy modifications and extensions to the solvers and turbulence models are possible. One such useful extension is the development of a projection and smoothing operation (see Refs. [3, 13]) that allows the use of experimental reference data [5]. Reference [5] used DAFoam and compared the effects of input data resolution on the assimilation accuracy by constructing synthetic data sets from DNS. Two-dimensional constraints could be correctly applied here due to the nature of the reference data used. Three-dimensional constraints have been applied using DAFoam to assimilate a 3D flow past a wall-mounted cube with a single plane of measurements generated using a hybrid RANS/LES model, demonstrating its feasibility in 3D assimilation, albeit with all three velocity components [20]. The operations defined in OpenFOAM and the ease of running applications on high-performance computing (HPC) clusters make DAFoam an attractive choice for performing 3D DA.

The complexity of performing 3D DA extends beyond traditional variational methods. A more recent technique, physics-informed neural networks (PINNs) [21], integrates the governing equations of fluid flow into the training of neural networks. Unlike traditional neural networks, PINNs enforce these equations as soft constraints by minimizing the residuals of the governing equations, ensuring that the model adheres to physical laws while learning the

solution. PINNs have been applied to fluid mechanics problems ranging from simple problems to demonstrate the method [22] to problems that have integrated turbulence models into the network [23]. Most of these studies use DNS or LES reference data to train the model. Some studies have explored the use of experimental data for jets [24] magnetic resonance imaging to correct noisy data and displacement artifacts [25], and a more complicated case of stalled flow past an airfoil at high Reynolds number [26]. The challenge of enforcing a 2D constraint for 3D data may not be as severe in PINNs since the constraints are not enforced as strongly as in the case of the variational method.

Most DA studies reviewed so far have focused on reconstructing flow quantities for canonical 2D problems. Despite the large-scale benefits of variational DA applied to steady-state cases, it is a technique that is undergoing development with limitations in assimilating 3D flows with limited observational data. This is reflected in the range of Reynolds numbers investigated (with the exception of Ref. [6]) and the primary use of high-fidelity reference data such as DNS/LES (with the exceptions of Refs. [6, 13]). The majority of corrections applied in these studies have targeted the turbulence model [4, 5, 20, 27], which has shown success for flows with moderate separation. When using high-fidelity simulations as reference data, the continuity constraint is strongly enforced, eliminating divergence errors. For weakly separated flows, span-averaging ensures a quasi-two-dimensional behavior, making the application of 2D continuity constraints reasonable during assimilation.

However, for high Reynolds number flows with strong separation and complex dynamics, access to high-fidelity data that strictly satisfies 2D continuity in a span-averaged sense is limited. In such cases, planar PIV reference data are often used. Experimental PIV data not only contains noise and uncertainties but also does not inherently satisfy 2D continuity. Performing 2D DA on such data while enforcing 2D continuity constraints leads to misinterpretation of the corrective/control variables, as briefly noted in Ref. [13]. This misinterpretation is problematic for developing generalizable models (as in Ref. [28]) and for deriving accurate second-order statistics, such as Reynolds shear stress, from assimilated mean velocity fields. When dealing with flows where the out-of-plane velocity component is significant, a 3D assimilation approach is necessary to accurately recover flow physics and improve the interpretation of the control variables. We present a novel approach for assimilating planar PIV data of a flow past a NACA0012 airfoil in deep stall at Reynolds number $Re_c \approx 7.5 \times 10^4$ [29] using variational data assimilation (DA) with a discrete adjoint method.

This method introduces a corrective forcing term in the momentum equations as a control variable while enforcing 3D constraints. We compare its performance against conventional variational DA with 2D constraints and re-evaluate the role of the corrective forcing and turbulence model.

The rest of the paper is organized as follows. The governing equations of fluid flow and the variational DA methodology are described in Sec. II, where we provide a comprehensive explanation of the assimilation framework. This section also describes the computational setup of 2D DA and introduces our novel method of 3D DA that is used in the study. In Sec. III, we present a NACA0012 airfoil in deep stall and discuss the experimental reference data set. The results from baseline computations are reported in Sec. IV. We compare the assimilation of experimental data using 2D constraints in Sec. V and 3D constraints in Sec. VI. The accuracy and impact on the reconstructed flow variables are analyzed in Sec. VII, with concluding remarks summarized in Sec. VIII.

II. TURBULENCE MODELING AND DATA ASSIMILATION

In this section, a discussion of the fluid model is provided along with the DA method used. Section II A presents the governing equations of fluid flow. Section II B provides a detailed description of the variational formulation and proceeds to motivate the choice of the control variable, where a forcing term is added to the momentum equations. In Sec. II C, we describe the methodology of 3D DA and distinguish it from a traditional 2D DA procedure.

A. Reynolds-Averaged Navier–Stokes

We are interested in the equations governing the ensemble-averaged flow. To this end, the ensemble-averaged (RANS) equations for an incompressible fluid are given by,

$$\frac{\partial U_i}{\partial x_i} = 0, \quad (1)$$

$$U_j \frac{\partial U_i}{\partial x_j} = -\frac{1}{\rho} \frac{\partial P^*}{\partial x_i} + \frac{\partial}{\partial x_j} \left(\nu \frac{\partial U_i}{\partial x_j} \right) - \frac{\partial \tau_{ij}}{\partial x_j}, \quad (2)$$

where U_i and P^* are the mean velocity components and pressure, respectively, ρ is the density of the fluid, ν is the kinematic viscosity, τ_{ij} is the Reynolds stress tensor, and x_i represents

the spatial coordinates. The Reynolds stress tensor $\tau_{ij} = \overline{u'_i u'_j}$ is defined as the averaged outer product of the fluctuating velocity components, presenting the well-known problem of closure. Here, the notation $\langle \cdot \rangle'$ denotes fluctuations, representing the deviation of the instantaneous velocity from its ensemble average, expressed as

$$u'_i = u_i - \bar{U}_i, \quad (3)$$

where u_i is the instantaneous velocity, and \bar{U}_i is the ensemble average. The overbar $\langle \cdot \rangle$ denotes Reynolds averaging.

B. Data assimilation methodology

We expect that the solution to the RANS equations, as defined in Sec. II A, will deviate from experimental observations, necessitating the use of DA. To apply variational DA, we require the definition of a control variable and an objective function. Following the approach of Ref. [2], we use a forcing term added to the momentum equations as the control variable. The modified momentum equations, incorporating this forcing term, are given by

$$U_j \frac{\partial U_i}{\partial x_j} + \frac{1}{\rho} \frac{\partial P^*}{\partial x_i} - \frac{\partial}{\partial x_j} \left(\nu \frac{\partial U_i}{\partial x_j} \right) - f_{R_i} - f_{c_i} = 0, \quad (4)$$

where f_{c_i} is the control variable applied as a corrective forcing in addition to $f_{R_i} = -\partial \tau_{ij} / \partial x_j$, the modeled Reynolds forcing. We use the Spalart-Allmaras [30] (SA) turbulence model (which will be referred to as the baseline model) to model the Reynolds stress tensor τ_{ij} . It applies the Boussinesq approximation through the eddy viscosity, which is obtained by solving the transport equation for a surrogate variable $\tilde{\nu}$. The functional form of the turbulence transport equation is given by:

$$\frac{D\tilde{\nu}}{Dt} = P(\tilde{\nu}, \mathbf{w}) + T(\tilde{\nabla} \nu, \mathbf{w}) - D(\tilde{\nu}, \mathbf{w}), \quad (5)$$

where \mathbf{w} is the vector of state variables such as mean velocity, pressure, and momentum flux, and P , T , and D are the production, diffusive transport, and destruction terms, respectively.

The corrective forcing term f_{c_i} is obtained by performing data assimilation. Using the Helmholtz decomposition, the corrective forcing term f_{c_i} can be decomposed into the sum of a solenoidal part f_{s_i} and an irrotational part given by,

$$f_{c_i} = f_{s_i} + \frac{1}{\rho} \frac{\partial \phi}{\partial x_i}. \quad (6)$$

The irrotational part is lumped in with the pressure $p^* \approx p - \phi$. The consequence of such a decomposition is that the recovery of the irrotational term is not possible unless pressure data are supplied.

The objective function is the discrepancy between high-fidelity/experimental and RANS simulation mean velocity fields,

$$J(\mathbf{u}, \mathbf{f}_c) = \frac{1}{2} \|\mathcal{Q}(\mathbf{u}, \mathbf{f}_c) - \tilde{\mathbf{Q}}\|_Q \quad (7)$$

where $\tilde{\mathbf{Q}}$ is a set of high-fidelity data/experimental measurements and bold-face symbols denote vectors. Operator $\mathcal{Q}(\cdot)$ extracts the computational data in such a way that $\mathcal{Q}(\mathbf{u}) \in Q$ is a projection of the computational mean velocity to the measurement space Q . $\|\cdot\|_Q$ is the generic norm in the measurement space. It should be noted that while the objective function J is computed only within the region where high-fidelity data or experimental measurements are available, the control variable f_c is defined over the entire computational domain. Consequently, it possesses the same number of degrees of freedom as the mesh size.

Variational DA is now formulated as an optimization problem where the goal is to minimize an objective function subject to some constraints. This is mathematically written as

$$\min_{\mathbf{f}_c \in \mathbb{R}^{n_{f_c}}} J(\mathbf{f}_c), \quad (8)$$

$$\text{s.t. } R(\mathbf{w}, \mathbf{f}_c) = 0, \quad (9)$$

$$\mathbf{f}_{c_L} \leq \mathbf{f}_c \leq \mathbf{f}_{c_U}, \quad (10)$$

where n_{f_c} is the size of the control variable, R is the governing equations that serve as constraints and \mathbf{f}_{c_L} and \mathbf{f}_{c_U} denote the lower and upper bounds, respectively, for the control variable. The bounds \mathbf{f}_{c_L} and \mathbf{f}_{c_U} are determined through iterative tuning to balance optimizer flexibility and solution stability. Starting from tight bounds near zero, the range is gradually widened if convergence issues arise. Excessively narrow bounds can restrict the optimizer, while overly loose bounds may slow convergence or destabilize the primal solver by allowing unrealistically large control values. For our test case, R represents the residual of the NS equations. In this formulation, the optimization is performed over the control variable \mathbf{f}_c , while the state variable \mathbf{w} is treated as an implicit function of \mathbf{f}_c and is obtained by solving the governing equations at each iteration. Equations 8 - 10 form a non-linear,

constrained minimization problem with equality and bound constraints and can be solved using gradient-based techniques.

Gradient-based optimization techniques require the total derivative of the objective function with respect to the control variable (hereafter referred to as sensitivity). An efficient way to compute the sensitivity is by employing an adjoint method, which ensures that the computational cost remains independent of the number of control variables [31]. We use the discrete adjoint method in this study for computing the sensitivities. If J and R are a univariate representation of the objective function and residual of the NS equations, respectively, the sensitivity can be computed using

$$\frac{dJ}{d\mathbf{f}_c} = \frac{\partial J}{\partial \mathbf{f}_c} - \psi^T \frac{\partial \mathbf{R}}{\partial \mathbf{f}_c}, \quad (11)$$

where ψ^T is the transpose of the adjoint vector. The detailed derivation can be found in Ref. [12]. DAFoam is used to obtain the sensitivity. DAFoam's source code is enriched with AD-forward (ADF) and AD-reverse (ADR) implementations using CoDiPack [32], enabling machine-precision gradient accuracy [12]. Once the sensitivity is obtained, we use sequential quadratic programming (SQP) [33] to perform the optimization. The SQP method solves an optimization problem by obtaining search directions from a sequence of quadratic programming (QP) sub-problems with linearized constraints. The sub-problem defines a quadratic model of a certain Lagrangian with linear equality and inequality constraints. A quasi-Newton method is used along with a line search filter that minimizes a penalty function to obtain search directions for the QP sub-problems. The initial set of iterates is then updated based on these search directions. We use SNOPT (Sparse Non-linear Optimizer) [Ref. [34]], which exploits the sparsity in the constraint Jacobian and employs a limited-memory BFGS (L-BFGS) approximation of the Hessian, allowing efficient handling of large-scale problems and addressing the issues mentioned earlier. The stopping criteria for the optimization are governed by the default convergence logic of the SNOPT optimizer, which monitors a combination of projected gradient norms and step sizes to determine when the optimization can no longer make meaningful progress.

We implement a cell-volume weighted averaging operator \mathcal{Q} that was used in Ref. [13] to project the computational mean velocity data onto the experimental grid. While this may be trivially accomplished using standard interpolation methods for high-fidelity reference data, the same cannot be employed for noisy experimental data. The cell-weighted

averaging procedure used in this study does not exploit cell-cell intersections for computing the volume-averaged mean velocity field as is done in Ref. [5]. Once the mean velocity data from the computational grid is projected onto the experimental grid, the objective function is calculated. Since the adjoint solution is forced on the computational grid, a smoothing operator \hat{Q} is used to transfer the discrepancy back to the computational grid. The implementation of this projection and smoothing operation is essential to ease the cost of calculating the discrepancy field for 3D problems. A custom objective function is also implemented that can utilize the projection and smoothing operations. The details of the implementation are presented in Appendix A.

C. Data assimilation setup

Two-dimensional DA (hereafter referred to as 2DVar) applies the governing equations in Eq. 2 and the constraints in Eq. 9 in only two spatial dimensions: streamwise and wall-normal. This approach has been the standard in all studies utilizing variational DA so far. However, for a truly 3D flow, enforcing 2D constraints leads to a misrepresentation of both the control variable and the role of the turbulence model. In this case, the control variable does not only compensate for the limitations of the turbulence model but also for the missing 3D effects, leading to an inaccurate correction. To address this, we propose a novel approach - 3D DA (hereafter referred to as 3DVar)—where the governing equations and constraints incorporate all three spatial dimensions: streamwise, wall-normal, and spanwise. This ensures that the 3D nature of the flow is properly accounted for, leading to a more physically meaningful interpretation of the control variable. We present the setups of both 2DVar and 3DVar and emphasize that, regardless of the chosen setup, the reference data always consist of a two-component mean velocity field.

Figure 1 illustrates a generic setup for 2DVar and 3DVar. The computational domain is assumed to be significantly larger than the experimental field of view, with L_x and L_y defining its extent. The computational grid may be block-structured, structured, or unstructured. The experimental data points are arranged uniformly in the two spatial directions, consistent with measurements from techniques such as PIV. In 2DVar, the spanwise length L_z is inconsequential as long as the number of points in the spanwise direction is $N_z = 1$. This is evident from the top view in Fig. 1, where the cell centroids from the experimental

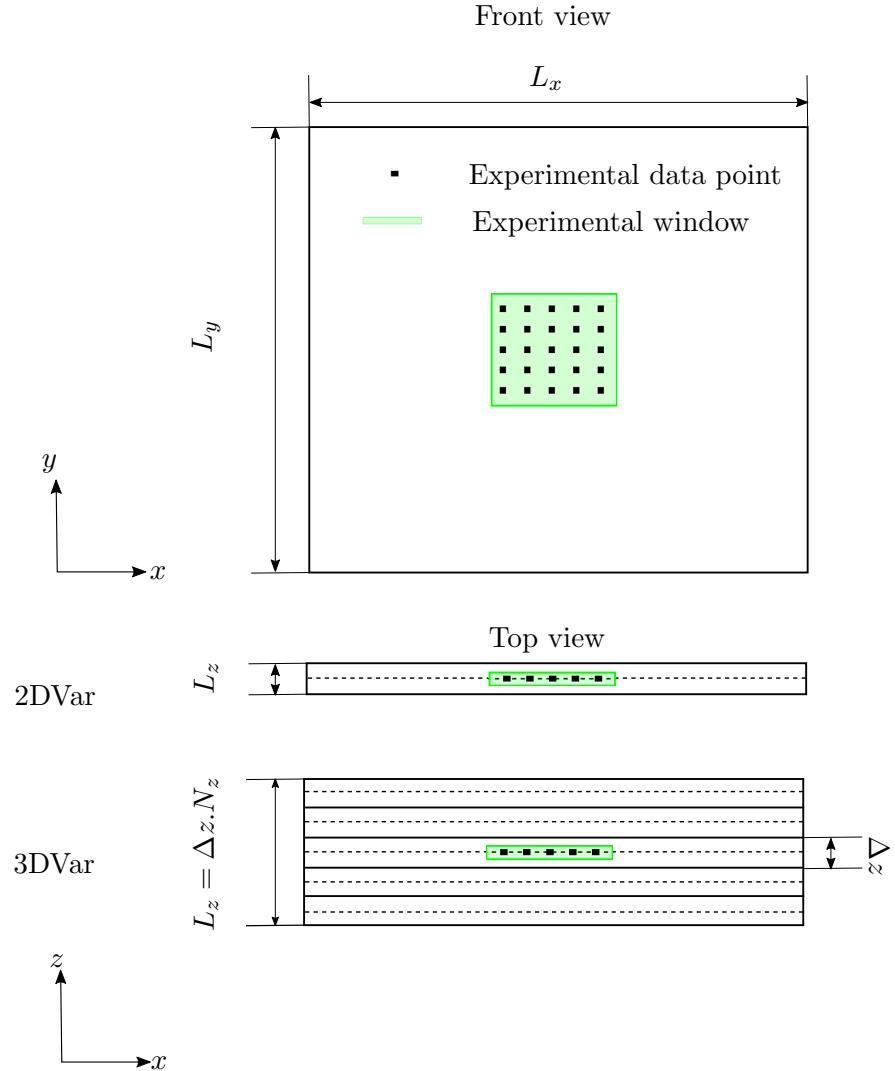


FIG. 1. Schematic of the 2DVar and 3DVar setups, showing front and top views of the computational domain along with the experimental field of view and data points. Here, L_x , L_y , and L_z are the streamwise, wall-normal, and spanwise domain lengths, respectively. The number of points in the spanwise direction is N_z , and the spanwise spacing of the computational mesh is Δz . Solid lines in the top view indicate a layer of cells, while dashed lines denote the position of the cell centroids.

data align with the computational grid, as indicated by the dashed lines.

The 3DVar setup differs in that the spanwise domain now extends by an integral multiple of the spanwise spacing Δz , given by $L_z = \Delta z \cdot N_z$, where $N_z > 1$. This effectively increases the number of divisions along the span for the same thickness. In Fig. 1, this is illustrated in the top view under the 3DVar section, where only one layer of cell centroids aligns with

the experimental data points. The spanwise spacing Δz is chosen to be slightly larger than the laser sheet thickness in a PIV setup and is defined by the limits z_{max} and z_{min} such that $\Delta z = z_{max} - z_{min}$, which is necessary for projection and smoothing operations described in Appendix. A. In the schematic shown (Fig. 1), $N_z = 5$; however, for practical cases, the choice of N_z depends on the specific problem. The spanwise domain must be sufficiently resolved to allow the development of the out-of-plane velocity component. A key point of similarity between both 2DVar and 3DVar is that the control variable f_c is optimized throughout the entire computational domain, allowing corrections even in regions where no data are available. The benefit of doing so, especially in the context of 3DVar will be highlighted later.

III. FLOW CONFIGURATION AND REFERENCE DATA

In this section, details about flow over an airfoil in deep stall are provided. Although the experimental data are known in a single plane, the flow is highly three-dimensional, which motivates a comparison between 2DVar and 3DVar. Section III A describes an airfoil in deep stall. The experimental data are described in Sec. III B, where we also highlight the divergence errors present in the mean velocity field.

A. Flow configuration

While we are interested in highly separated flows in general, we aim to demonstrate our methodology on a specific case of airfoil stall. The reason for this choice is the complex physics and flow features that are involved in the process of airfoil stall, including, but not limited to, separation (leading or trailing edge), presence of leading and trailing edge shear layers, and recirculating flow. Additionally, such strong streamline curvature and the presence of a pressure gradient make it challenging for existing turbulence models to capture the flow physics. We use a NACA0012 airfoil at an angle of attack $\alpha = 15^\circ$ and a chord-based Reynolds number $Re_c \approx 7.5 \times 10^4$, which is in the deep stall regime.

B. Experimental reference data

The reference mean velocity fields are obtained from the experiment performed in Ref. [29]. Time-resolved, planar PIV measurements are taken in a water flume facility. The test section is 6.75 m long, 1.2 m wide and 0.5 m deep. The airfoil has a chord of 15 cm and a span of 70 cm. More than half the span is submerged. Three 4-megapixel, high-speed Phantom Veo 640-S cameras are used to capture the images. The forces are measured using a six-axis force/torque load cell (ATI Delta IP65) mounted to the airfoil and synchronized to the PIV acquisition using an NI USB-6251 data acquisition device (DAQ). The acquisition frequency is 1kHz. Multipass PIV is performed using a verified in-house MATLAB code with 3 passes per window size and square windows decreasing from 64 by 64 pixels to 32 by 32 pixels to 24 by 24 pixels with 50% overlap. This results in 197×249 vectors with 181 vectors along the airfoil chord. Further details can be found in Ref. [29]. A generous mask up to $y^+ \approx 100$ is used, removing the data close to the airfoil surface that is corrupted from surface reflections.

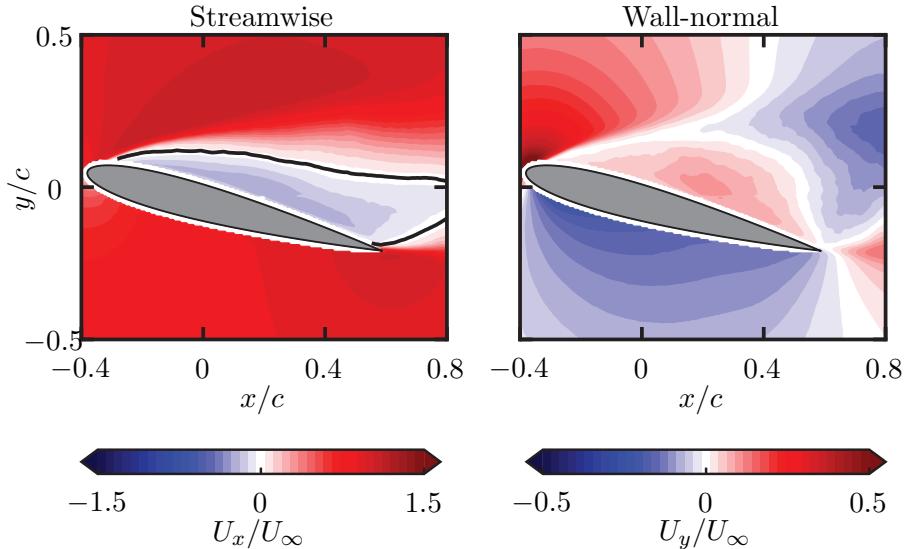


FIG. 2. Mean streamwise and wall-normal velocity fields obtained from experiment. The $U_x = 0$ contour is shown as solid black line.

The experimental mean streamwise and wall-normal velocity fields are shown in Fig. 2. We plot the $U_x = 0$ contour to delineate the region of reverse flow. The mean streamwise velocity field is representative of a deep stall case with leading edge separation and a large recirculation bubble. This is a case of a well-resolved flow field but the field of view (FOV)

is limited. The streamwise extent of the FOV is about $1.2c$ while the wall-normal extent is $1c$. The full extent of the recirculation bubble is not captured.

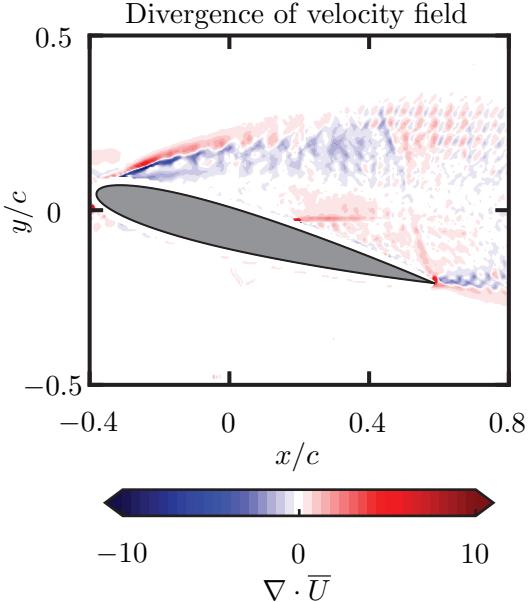


FIG. 3. The divergence of the experimental reference data.

A key aspect of the experimental dataset is the presence of divergence errors. We only have access to the streamwise and wall-normal components of velocity in this experimental dataset. As a result, the divergence of the velocity field would lack the contribution of the out-of-plane velocity gradient (in this case, $\partial U_z / \partial z$). The 2D divergence field of this experimental dataset can then be calculated using the streamwise and wall-normal velocity gradients only as

$$\nabla \cdot U = \partial U_x / \partial x + \partial U_y / \partial y. \quad (12)$$

Since the data are present on a uniform grid, a central differencing scheme (accurate up to second order) is used to compute the 2D divergence, which is presented in Fig. 3. There are large values in the leading edge shear layer. Additionally, there is an unphysical artifact from the trailing edge of the airfoil that is caused due to the overlap of the laser sheets from the pressure and suction sides of the airfoil. The divergence in the wake of the airfoil appears grainy and coincides with the region of high experimental uncertainty. This is expected, as the wake in deep stall is dominated by large-scale unsteady vortex shedding and elevated turbulence levels, which increase the root mean square of the velocity fluctuations. Since random uncertainty scales with $\sqrt{u'^2}$ and $\sqrt{v'^2}$ (where u' and v' are fluctuating streamwise

and wall-normal velocity components, respectively), the standard error ϵ_r naturally rises in this region. Additionally, the unsteady nature of the flow can lead to correlation peak broadening and loss of particle pairs, further elevating measurement uncertainty [35, 36]. All these features are quite typical of an experimental dataset. The 2D divergence of the dataset on a discrete grid is then calculated as,

$$\nabla \cdot U = \frac{1}{N_e} \sum_{N_e} \left(\frac{\partial U_x}{\partial x} + \frac{\partial U_y}{\partial y} \right). \quad (13)$$

The magnitude of the 2D divergence is $\nabla \cdot U = -0.0359$. The use of this dataset for 2DVar assumes that the velocity field obeys 2D continuity. It is evident from Fig. 3 and the magnitude of 2D divergence that this is not the case. In reality, the presence of noise would make it challenging to achieve a truly divergence-free velocity field. However, such a large discrepancy could only be explained by the three-dimensionality of the flow field. This would require the contribution of the out-of-plane velocity gradient to have a fully divergence-free velocity field within experimental uncertainty.

Such three-dimensionality is not unexpected in turbulent separated flows at moderate Reynolds numbers, particularly in the deep stall regime. Prior studies have reported features such as spanwise vortex structures and bluff-body-like shedding in similar configurations [37, 38]. While a slight misalignment of the laser sheet is present in the experimental setup [29], its influence appears mostly confined to the wake. In contrast, the leading-edge region, where divergence errors are largest, is less affected by this artifact. Here, the inferred transverse gradients, supported by assimilation results, suggest the presence of spanwise flow potentially linked to coherent structures like stall cells. Although such features are unsteady and may be obscured by time-averaging, they remain a plausible source of the observed divergence.

IV. BASELINE COMPUTATION

With the reference datasets defined and characterized in Sec. III, we present the details of the baseline simulations. The choice of the computational mesh and boundary conditions for 2DVar and 3DVar assimilation are discussed in Sec. IV A. The results from the baseline computation are then compared with the reference data in Sec. IV B.

A. Setup for baseline computation

We use the SA model as the baseline for both 2DVar and 3DVar, as reported in Sec. II B. The RANS equations are solved using the Semi-Implicit Method for Pressure Linked Equations (SIMPLE) [39] through OpenFOAM’s inbuilt `simpleFOAM` solver. The gradients are calculated using a second-order accurate central differencing scheme. The flow variables (velocity and $\tilde{\nu}$) are discretized using a first-order upwind method to promote the stability of the solver. Each matrix equation is solved using the Gauss-Seidel method. Convergence is determined based on a tolerance of 10^{-5} for the residual of pressure and velocity components. We study two separate cases: 2DVar assimilation of experimental reference data, and 3DVar assimilation of experimental data. Regardless of the case, all baseline computations use the SA model by setting the corrective forcing term $f_{c_i} = \mathbf{0}$ in Eq. 4. By setting the forcing term f_{c_i} to zero, it means that the initial guess of the control variable for both 2DVar and 3DVar assimilation is zero. We set $U_\infty = 1 \text{ ms}^{-1}$, $c = 1 \text{ m}$ and $\nu = 1/Re = 1.333 \times 10^{-5} \text{ m}^2 \text{s}^{-1}$.

The grid used for 2DVar is a C-type grid shown in Fig. 4. The domain extends by $16c$ downstream and $8c$ upstream from the leading edge of the airfoil and $0.1c$ in the out-of-plane direction. Unless otherwise specified, the positive x is designated as streamwise, the positive y as wall-normal, and the positive z as spanwise directions and is part of the right-handed coordinate system. The computational mesh consists of mostly hexahedral cells and has a refinement in the region close to the airfoil to ensure that the shear layers and wake are well resolved. Inflation layers are added to resolve the viscous sub-layer so that wall functions are not needed. The 2D grid is extruded by a single layer of cells in the spanwise direction such that $N_z = 1$ where N_z is the number of points in the spanwise direction. This is required since all problems are inherently treated as 3D in OpenFOAM. We will solve the problem as 2D by enforcing the relevant boundary conditions. The discretization results in 400 points along the airfoil and a total of 44687 cells in the entire domain.

The boundary conditions for velocity, pressure, eddy viscosity ν_t and modified eddy viscosity $\tilde{\nu}$ are summarized in Tab. I. We employ a standard velocity inlet, pressure outlet boundary condition for incompressible flows where the velocity and pressure are set to fixed values at the inlet and outlet of the domain, respectively. We set a fixed value inlet boundary for $\tilde{\nu}_\infty = 3\nu_\infty$ to $5\nu_\infty$ with the recommendation from Ref. [30] and a fixed gradient outlet

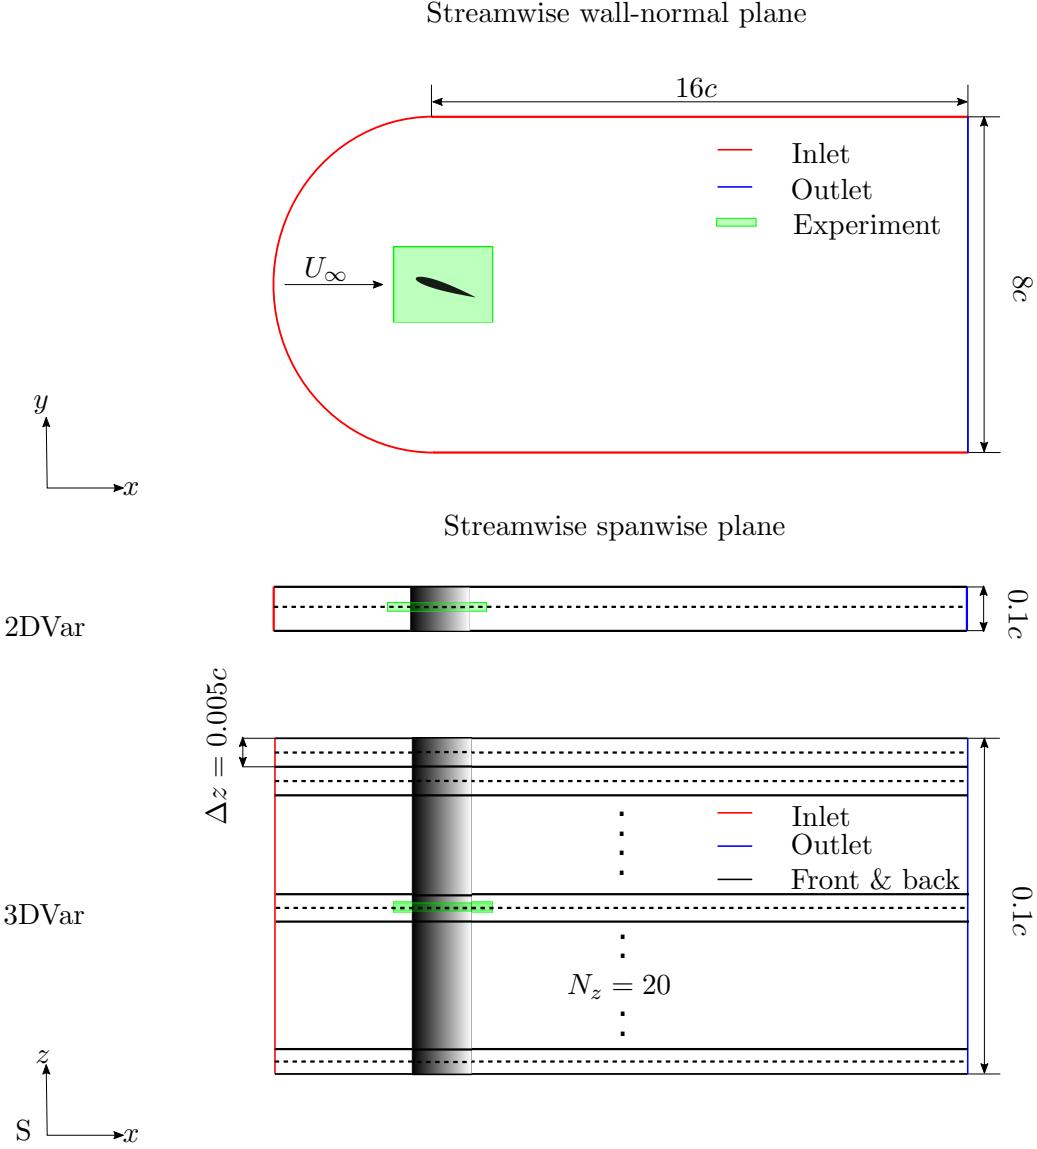


FIG. 4. The computational domains for 2DVar and 3DVar are shown, highlighting the streamwise-wall-normal and streamwise-spanwise planes. For clarity, the spanwise length $L_z = 0.1c$ has been exaggerated in the 3DVar representation to better illustrate the divisions. The inlet, outlet, and lateral faces of the domain are labeled, along with the flow direction. The experimental field of view is also indicated, and the cell centroids are represented by dashed lines.

boundary. The lateral faces of the domain are designated as a no-flow condition. The presence of inflation layers and the resolution of the viscous sub-layer means that there are no wall functions required, setting a value of $\tilde{\nu} = 0$ at the wall. Since the transport equations are solved for $\tilde{\nu}$ as outlined in Eq. 5, the boundary condition for the eddy viscosity used to calculate the closure term is designated as calculated in OpenFOAM that estimates the

TABLE I. Boundary conditions for baseline SA model

	Inlet	Outlet	Front & back	Airfoil
U	1	$\frac{\partial U}{\partial n} = 0$	$\hat{n} \cdot \vec{U} = 0$ & $\frac{\partial U}{\partial n} = 0$	no slip
p	$\frac{\partial p}{\partial n} = 0$	0	$\frac{\partial p}{\partial n} = 0$	$\frac{\partial p}{\partial n} = 0$
$\tilde{\nu}$	$\tilde{\nu}_\infty$	$\frac{\partial \tilde{\nu}}{\partial n} = 0$	$\frac{\partial \tilde{\nu}}{\partial n} = 0$	$\tilde{\nu} = 0$
ν_t	calculated	calculated	$\frac{\partial \nu_t}{\partial n} = 0$	calculated

value of ν_t based on a relation provided in Ref. [30].

The mesh used for 3DVar is different from that used for 2DVar. As described in Sec. II C, the computational grid for 3DVar can be generated by increasing the number of divisions along the span to $N_z = 20$ with the thickness of each layer in the spanwise direction being $\Delta z = 0.005c$. This results in a computational mesh size of 893,740 cells. A schematic of this setup, showing the spanwise resolution, can be seen in Fig. 4. The boundary conditions are the same as those summarized in Tab. I. The only difference is that the spanwise direction has sufficient resolution to allow the development of an out-of-plane velocity component that allows the enforcement of 3D continuity. The thickness of each spanwise layer is chosen to match the laser sheet thickness used in the PIV experiment, ensuring that large gradients are not poorly approximated due to insufficient discretization. An important consideration is the overall spanwise extent of the computational domain. In this study, we restrict ourselves to assimilating a small volume surrounding the experimental data plane, rather than the full wingspan used in the experiment. The goal is not to reproduce the global three-dimensional features of the flow field, but to focus on the local three-dimensionality captured within the measurement plane. We speculate that a larger computational domain could allow the capture of spanwise-organized flow structures, such as stall cells. However, this is beyond the scope of the present study.

The discrepancy between the baseline and reference mean velocity data will be quantified by using the L_1 norm field which is computed as

$$L_1 = \frac{|\mathcal{Q}(U_x) - \hat{U}_x| + |\mathcal{Q}(U_y) - \hat{U}_y|}{U_\infty}, \quad (14)$$

where \hat{U}_x , \hat{U}_y are the reference data mean streamwise and wall-normal velocities, respectively, and \mathcal{Q} is the projection operator which depends on the reference data set. This projection

operator allows the computation of the L_1 norm on the experimental grid.

B. Results from baseline computation

The comparison of the mean streamwise and wall-normal velocity field obtained from the baseline model with experimental reference data is presented in Fig. 5. We compare the extent of recirculation using the $U_x = 0$ contour where the baseline model is represented as a dashed line and the experimental data as a solid one, and also plot the L_1 norm. The experimental data are on a structured grid with a uniform distribution of points, whereas the baseline model is applied on a computational mesh that is unstructured. We use Eq. 14 to compute the L_1 norm where the operator \mathcal{Q} is linear interpolation to transfer the mean velocity data from the computational grid to the experimental one. The L_1 norm field reveals a very large discrepancy in the leading edge shear layer also corroborated by a reduced recirculation bubble observed in the baseline model. Such a large discrepancy is a clear indication that RANS turbulence models lack the accuracy required to model high Reynolds number airfoils in stall.

The baseline model used for 2DVar solves the governing equations in 2D. The baseline model results for 3DVar are not shown here, as the streamwise and wall-normal velocity components obtained from applying the baseline model on a 3D computational mesh do not differ significantly from those in 2DVar. The spanwise component of the mean velocity, computed using the baseline model, is negligibly small (on the order of 10^{-7}). While the computational mesh is sufficiently resolved along the span, limitations of the model itself may prevent accurate resolution of the spanwise velocity component.

V. 2DVAR ASSIMILATION OF EXPERIMENTAL DATA

We proceed to perform 2DVar using experimental reference data that do not satisfy 2D continuity computed using Eq. 12 as tabulated in Tab. V. The baseline model used is the SA RANS turbulence model. We use a mesh with a single cell in the spanwise direction, essentially treating the problem as 2D. The experimental reference data are limited to a much smaller domain owing to the reduced field of view obtained through PIV. The observation operator \mathcal{Q} in Eq. 7 projects the mean velocity data from the CFD domain within this limited

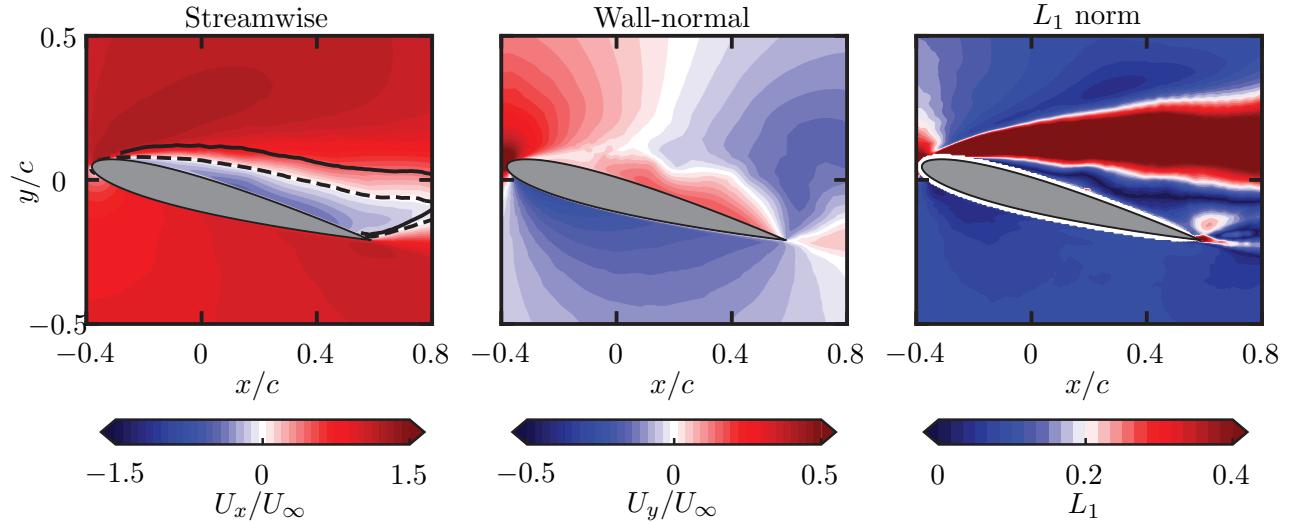


FIG. 5. From left to right: mean streamwise velocity, mean wall-normal velocity fields, and the L_1 norm comparing mean velocity fields from the baseline SA RANS turbulence model with the reference experimental data. The $U_x = 0$ contour is depicted as a solid black line for the experimental reference data and as a dashed black line for the baseline model in the mean streamwise velocity field.

field of view onto the experimental grid, and as such the objective function is calculated only within this region. We examine the mean velocity field and the control variable, i.e. the corrective forcing f_{c_i} and compare it to the Reynolds forcing f_{R_i} .

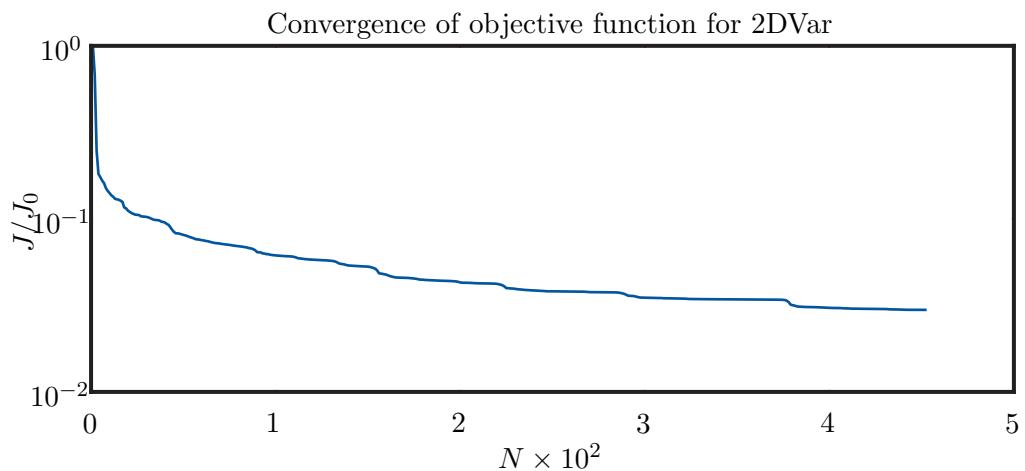


FIG. 6. Convergence of objective function J scaled by its value at the first optimization iteration J_0 using 2DVar for experimental reference data. The number of optimization iterations is given by N .

The convergence of the scaled objective function is shown in Fig. 6. The convergence is rapid in the initial iterations, followed by a more gradual change as the optimization progresses. This behavior is expected, as the discrepancy is large at the beginning, prompting the optimizer to apply significant forcing to reduce the discrepancy quickly, with subsequent iterations yielding smaller, more gradual improvements. The scaled objective function converges by over an order of magnitude, indicating successful assimilation of the mean velocity field. Convergence is achieved in just under 500 optimization iterations. The high iteration count is expected, as the 2D setup must compensate for both model-form errors and the lack of out-of-plane flow information, leading to a more constrained and ill-posed optimization landscape. The assimilation is performed using 12 cores of dual 2.0 GHz Intel Xeon Gold 6138 processors, resulting in 11.89 GB of memory usage. Given the parallel computing capabilities of DAFoam on HPC resources, the overall computational cost remains manageable. The assimilated mean streamwise and wall-normal velocity fields are presented in Fig. 7. The discrepancy field is quantified using the L_1 norm. The improvement in the mean velocity fields is evident from the reduced discrepancy in the leading-edge shear layer in Fig. 7, compared to the baseline. This reduction in discrepancy leads to an increase in the recirculation bubble size, as indicated by the improved agreement between the $U_x = 0$ contours of the experimental and data-assimilated fields.

To quantify the improvement in the prediction of mean velocity by assimilation, we define an integrated discrepancy which is calculated as,

$$\epsilon = \frac{1}{N_e} \sum_{i=1}^{N_e} \frac{|\mathcal{Q}(U_x^i) - \hat{U}_x^i| + |\mathcal{Q}(U_y^i) - \hat{U}_y^i|}{U_\infty}, \quad (15)$$

where N_e is the number of experimental grid points and the superscript i is the i^{th} grid point in the experiment. Equation 15 is just an integral of Eq. 14. This is used to compute the relative error RE as

$$RE = \frac{\epsilon_A}{\epsilon_B}, \quad (16)$$

where ϵ_A and ϵ_B are the integrated discrepancies of the assimilated and baseline mean velocity fields, respectively, compared to the reference data. Smaller values of RE show that the assimilation outperforms the baseline. The values for ϵ_B , ϵ_A , and RE are tabulated in Tab. II. We observe that the integrated discrepancy of the assimilated model is almost

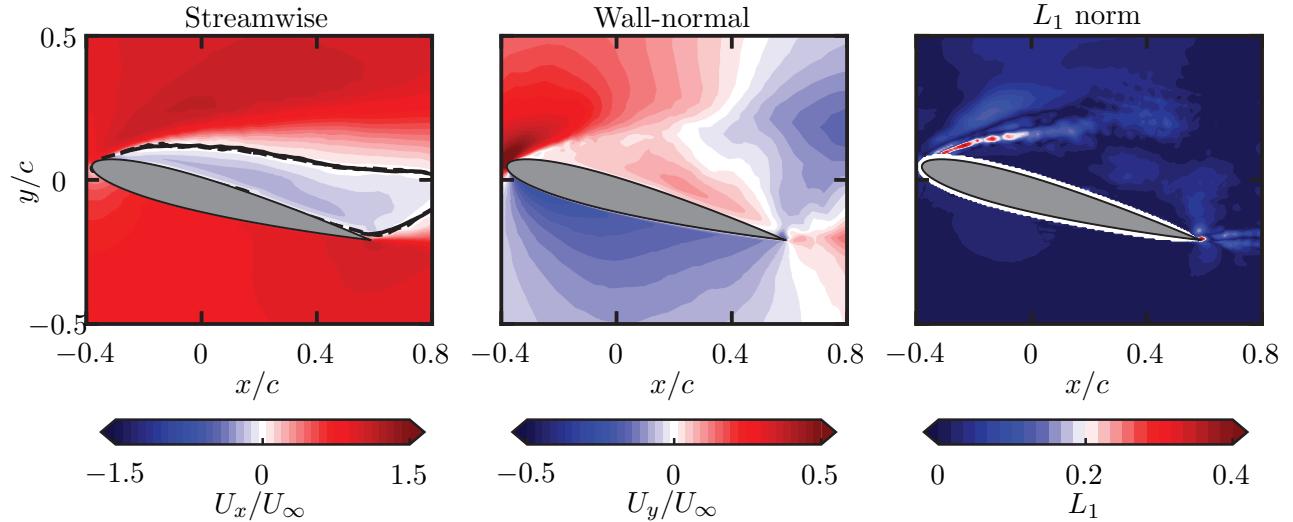


FIG. 7. From left to right: 2DVar-assimilated mean streamwise velocity, mean wall-normal velocity fields, and the L_1 norm comparing the assimilated mean velocity fields with the experimental data. The $U_x = 0$ contour is depicted as a solid black line for the experimental reference data and as a dashed black line for the data-assimilated model.

an order of magnitude smaller compared to the baseline which is also reflected in the low magnitude of the RE .

TABLE II. Integrated discrepancy and relative error for 2DVar

	ϵ_B	ϵ_A	RE
Magnitude	0.164	0.019	0.119

To evaluate the divergence of the data-assimilated field, we first compute the mean velocity gradients on the computational grid using the Green-Gauss method. These gradients are then interpolated onto the experimental grid following the procedure described in Appendix A, ensuring that conservation properties are preserved as much as possible. The divergence is subsequently calculated using Eq. 13 and is found to be $\nabla \cdot U \approx 5 \times 10^{-4}$. This value is two orders of magnitude smaller than the divergence of the experimental mean velocity field, as reported in Tab. V. However, discounting interpolation errors, the assimilated mean velocity field can still be considered approximately divergence-free. The discrepancy in divergence between the experimental and assimilated mean velocity fields arises from enforcing a strict 2D divergence-free constraint (Eq. 12) while assimilating a mean velocity field that does not

strictly satisfy such a 2D constraint.

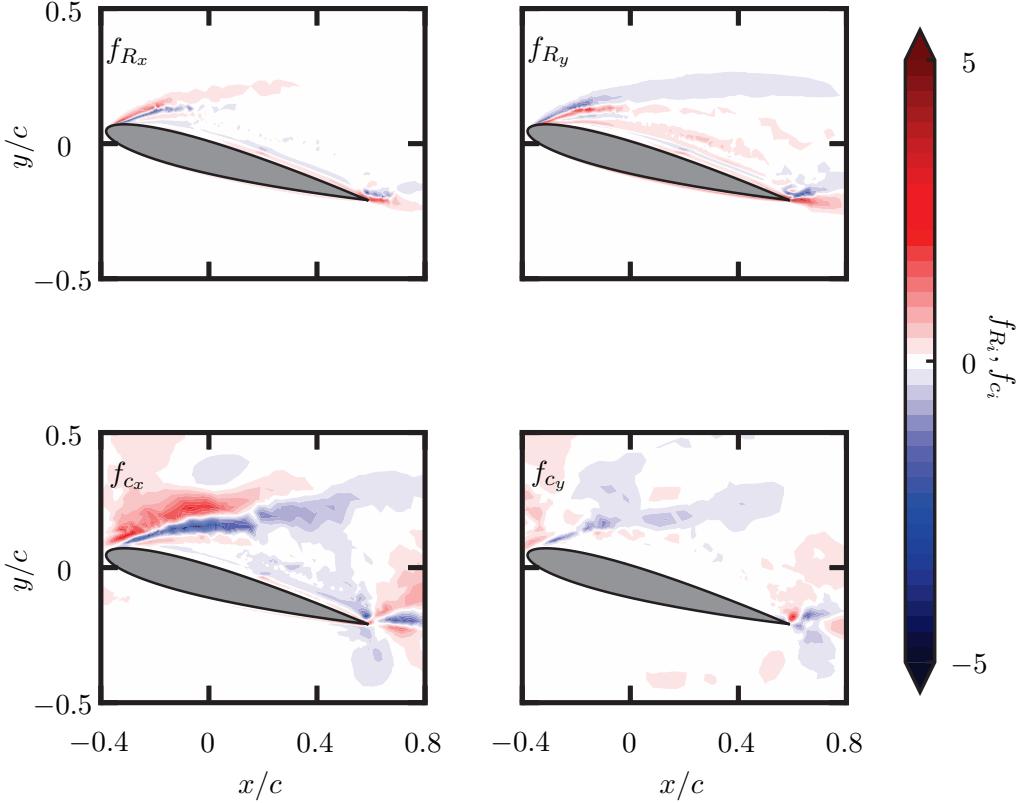


FIG. 8. Comparison of the streamwise and wall-normal components of the corrective forcing f_{c_i} with the Reynolds forcing f_{R_i} for the 2DVar assimilation using experimental reference data.

Such a discrepancy in the divergence between the experimental and 2DVar assimilated mean velocity fields directly impacts the corrective forcing field and its interpretation. A comparison between the corrective forcing and Reynolds forcing for the 2DVar assimilated field, using experimental reference data, is presented in Fig. 8. The Reynolds forcing is dominant in the leading- and trailing-edge shear layers. The streamwise component of the corrective forcing exhibits significant spread at the leading edge, with alternating regions of positive and negative magnitude tracing the shear layer. At first glance, it would be reasonable to expect that the forcing in the shear layer entrains freestream flow into the recirculation bubble, leading to its expansion. However, a closer examination reveals that the streamwise corrective forcing aligns with regions in the leading-edge shear layer where the divergence of the experimental velocity field is high, as seen in Fig. 3. This makes it challenging to attribute the effect of the forcing solely to compensating for the lack of physics captured by RANS turbulence models, as it is possible that the forcing is instead

TABLE III. Maximum and Minimum Values for Forcing Components

	Corrective Forcing		Reynolds Forcing	
	f_{c_x}	f_{c_y}	f_{R_x}	f_{R_y}
Minimum	-4.64	-1.60	-3.97	-9.6
Maximum	3.30	2.82	3.96	5.9

addressing the absence of a divergence-free velocity field.

To quantify these effects, we compare the maximum and minimum values of the corrective forcing with those of the Reynolds forcing in Tab. III. In general, the corrective forcing values are of comparable magnitude to the Reynolds forcing, except for the minimum value of the wall-normal component. A key factor contributing to this similarity is that the turbulence model cannot account for certain features present in experimental data, such as divergence errors and measurement noise. In this case, the experimental mean velocity field exhibits a significant divergence error, likely due to the inherently 3D nature of the flow, which the 2DVar assimilation is attempting to constrain within a 2D framework. These discrepancies are effectively absorbed into the forcing term. As a result, the corrective forcing not only captures unresolved physical effects in RANS models but also compensates for experimental artifacts. This blending makes it difficult to isolate its true contribution from setup-specific adjustments, posing challenges for symbolic regression, where models may inadvertently reflect experimental-specific features.

VI. 3DVAR ASSIMILATION OF EXPERIMENTAL DATA

We present the results of our novel approach of performing 3DVar assimilation using 2D2C experimental reference data. The key distinction between 3DVar and 2DVar lies in the nature of the constraints enforced. In 3DVar, 3D constraints are enforced through sufficient resolution in the spanwise direction. It will be shown that the enforcement of the 3D continuity constraint (as shown in Eq. 17) allows the development of an out-of-plane component of velocity, confirming the 3D nature of the flow and the inconsistencies in applying a 2D constraint as in 2DVar.

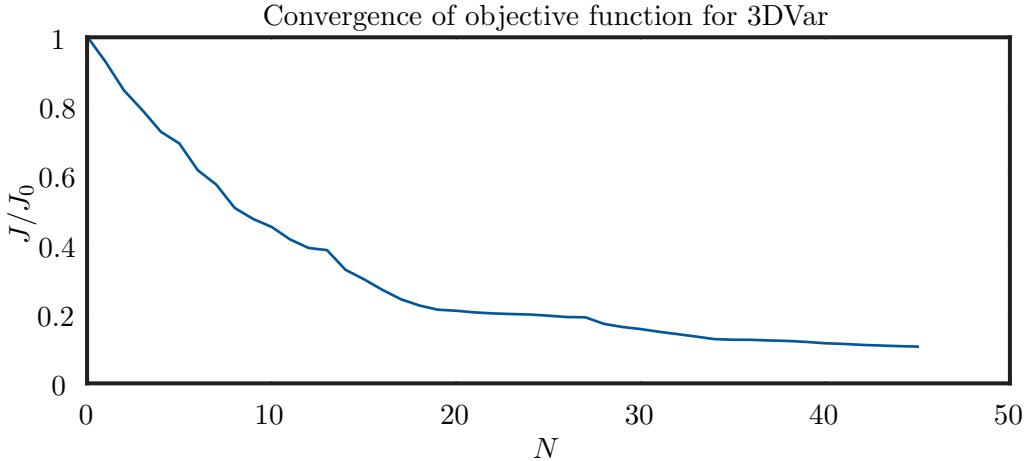


FIG. 9. Convergence of objective function J scaled by its value at the first optimization iteration J_0 using 3DVar for experimental reference data. The number of optimization iterations is given by N .

The convergence of the scaled objective function is presented in Fig. 9. The optimization converges in 45 iterations with a monotonically decreasing trend. Convergence required fewer iterations due to added spanwise degrees of freedom, which give the optimizer greater flexibility, especially in satisfying the divergence constraint, resulting in more efficient corrections. The convergence of the objective function is also a lot more gradual in the case of 3DVar as compared to 2DVar (see Fig. 6). The application of 3D constraints, in contrast to 2D constraints, may lead to a less aggressive behavior of the optimizer in adjusting the control variable, as will be discussed in the following section. The assimilation is carried out using 64 cores of dual 2.0 GHz Intel Xeon Gold 6130 processors on a high-memory node, resulting in a memory usage of 490 GB. Although solving adjoint equations using the discrete adjoint method typically demands significant memory, techniques like graph coloring in DAFoam make it feasible to handle this case efficiently.

The scaled objective function reduces almost by an order of magnitude, which indicates that the assimilation has taken place successfully. The assimilated mean streamwise and wall-normal velocity fields from 3DVar are presented in Fig. 10. The data are extracted from the assimilating plane at $z=0$. The discrepancy along the leading-edge shear layer is reduced, as indicated by the L_1 norm, and the 3DVar assimilation correctly predicts the recirculation bubble size. However, the discrepancy field suggests that 3D assimilation fields do not converge as closely to the experimental data as they do in 2DVar. This is expected

TABLE IV. Integrated discrepancy and relative error for 3DVar

	ϵ_B	ϵ_A	RE
Magnitude	0.164	0.061	0.371

since 3DVar, owing to its tighter enforcement of the governing equations, is less capable of absorbing experimental imperfections through unphysical corrections (as observed in the case of 2DVar) and thus avoids overfitting the data. This is reflected in the integrated discrepancy as tabulated in Tab. IV, where $\epsilon_A = 0.061$, which is a little over three times the integrated discrepancy of the 2DVar assimilated field. Some error remains in the leading-edge shear layer that has not been minimized. Despite this, the recirculation region in the assimilated field aligns well with the experimental data.

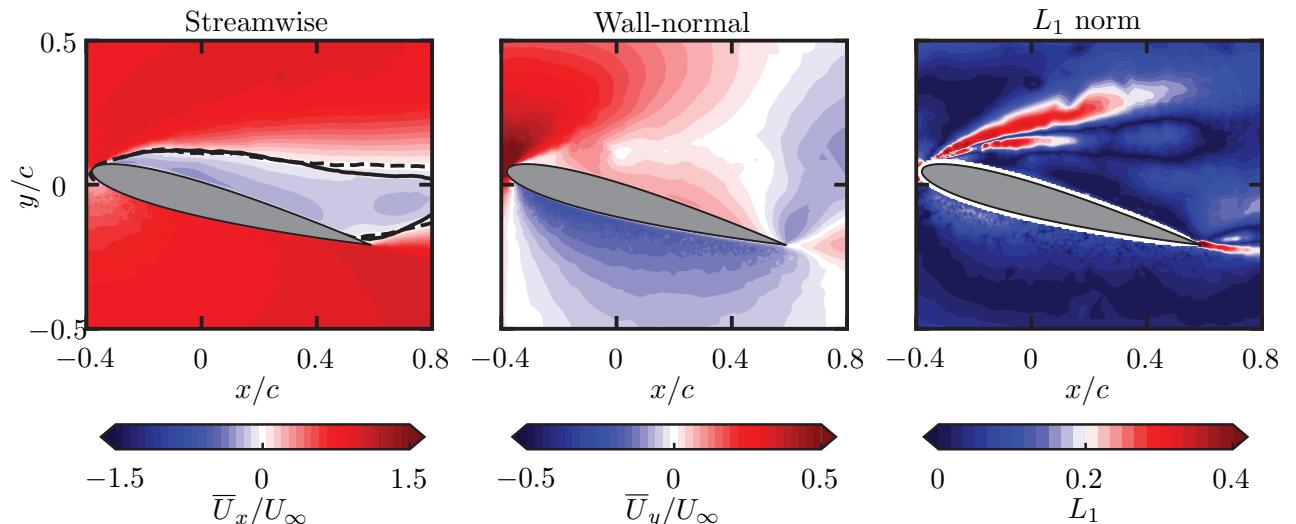


FIG. 10. From left to right: 3DVar-assimilated mean streamwise velocity, mean wall-normal velocity fields, and the L_1 norm comparing the assimilated mean velocity fields with the experimental data. The $U_x = 0$ contour is depicted as a solid black line for the experimental reference data and as a dashed black line for the data-assimilated model.

One inconsistency in 2DVar is the enforcement of a strict 2D divergence-free constraint, which ensures the assimilated mean velocity field satisfies this condition—even though the experimental reference data does not. In 2DVar, the divergence is computed using only two velocity gradients, as shown in Eq. 12. Here, we extend this analysis to the 3DVar assimilated mean velocity field. Following the same approach as in the 2DVar case, we

compute $\partial U_x / \partial x$, $\partial U_y / \partial y$ and $\partial U_z / \partial z$ on the computational mesh using the Green-Gauss method. The gradients along the $z = 0$ plane are then extracted and interpolated onto the experimental grid using volume weighting, as described in Appendix A. The divergence of the velocity field is then calculated as

$$\nabla \cdot U = \frac{1}{N_e} \sum_{N_e} \left(\frac{\partial U_x}{\partial x} + \frac{\partial U_y}{\partial y} + \frac{\partial U_z}{\partial z} \right). \quad (17)$$

We compute the integrated spanwise velocity gradient separately to lend weight to the assertion that flow is indeed 3D given by

$$\nabla_z U = \frac{1}{N_e} \sum_{N_e} \frac{\partial U_z}{\partial z}, \quad (18)$$

where $\nabla_z U$ denotes an integrated spanwise velocity gradient of the 3DVar assimilated field. To further assess the consistency between the 3DVar assimilated and experimental fields, we also compute a 2D divergence measure using only $\partial U_x / \partial x$ and $\partial U_y / \partial y$ by applying Eq. 12. This allows us to verify whether the 2D divergence of the 3DVar assimilated mean velocity field matches that of the experiment. The magnitudes of 2D and 3D divergence for the 3DVar assimilated mean velocity field, the divergence of the experimental data, and the integrated spanwise velocity gradient are tabulated in Tab. V.

TABLE V. Divergence values for 3DVar assimilation

	Experiment	2D Divergence	3D Divergence	$\nabla_z U$
Divergence	-0.0359	-0.0358	-4×10^{-4}	0.0354

There are a few important observations to highlight. First, the 2D divergence of the 3DVar assimilated mean velocity field aligns closely with that of the experimental mean velocity field (see Tab. V). Second, the 3D divergence of the assimilated field is of the same order of magnitude as the 2DVar assimilated field, with the added contribution from the spanwise velocity gradient. By extending the spanwise domain, we have effectively lifted the restrictive 2D continuity constraint, allowing for the development of a spanwise velocity component. The inclusion of this spanwise velocity gradient through the 3D constraints ensures that the 3DVar assimilated field is genuinely divergence-free, while still maintaining agreement in 2D divergence with the experiment. This demonstrates that the flow is inherently 3D, and attempting to assimilate it using 2D constraints is not appropriate.

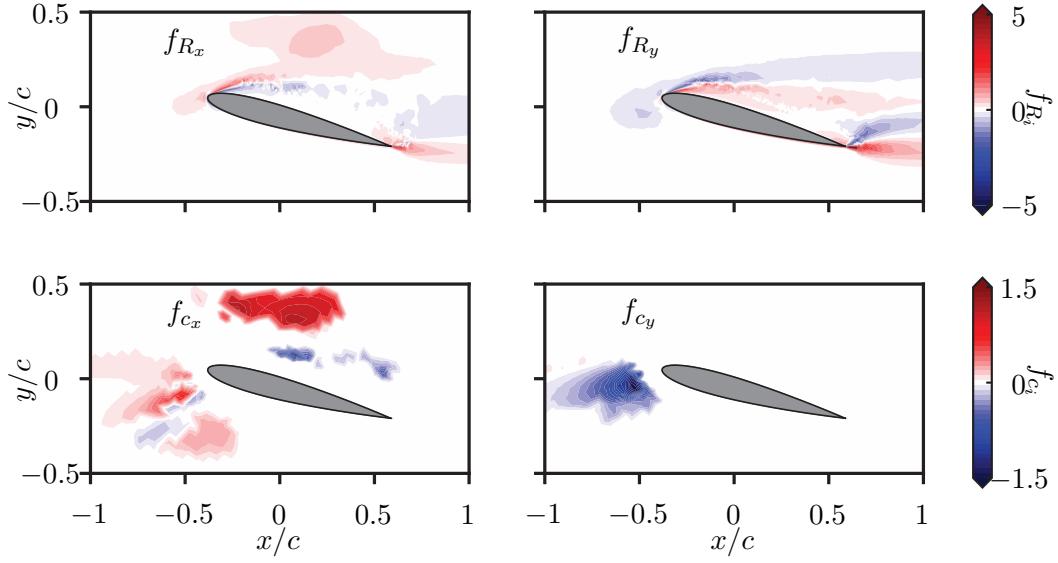


FIG. 11. Comparison of the streamwise and wall-normal components of the corrective forcing f_{c_i} with the Reynolds forcing f_{R_i} for the 3DVar assimilation using experimental reference data.

With half the divergence magnitude in 3DVar assimilation coming from $\nabla_z U$ (as computed in Eq. 18), we turn our attention to the corrective and Reynolds forcing fields. The comparison between these forcing terms is shown in Fig. 11. The streamwise component of the Reynolds forcing remains largely similar to that of 2DVar (Fig. 8), whereas the wall-normal Reynolds forcing exhibits a much broader spread in the leading and trailing edge shear layers. In contrast, the corrective forcing differs significantly from its 2DVar counterpart. Notably, there is some forcing just upstream of the leading edge, which at first glance may appear nonphysical. To interpret this corrective forcing, we consider several sources of uncertainty in the experimental setup, including freestream velocity, airfoil position, and angle of attack. The location and magnitude of the corrective forcing suggest that it is effectively slowing down the velocity at a given angle, potentially accounting for these uncertainties. Additionally, the level of freestream turbulence in the experiment does not match that of the assimilation, meaning this forcing term could also be compensating for that discrepancy. The presence of forcing outside the experimental field of view highlights that the control variable is optimized over the entire domain, enabling the assimilation to

compensate for discrepancies even in regions where no data are provided. The streamwise corrective forcing f_{C_x} shows clear spanwise variation, peaking along the data plane ($z/c = 0$) and diminishing off-plane, reflecting the localized influence of the available velocity data. This contrasts with f_{C_y} , which varies primarily in magnitude but remains spanwise-uniform, consistent with the weaker sensitivity of wall-normal velocity and the dominant contribution of f_{C_x} .

TABLE VI. Maximum and minimum values for forcing components in 3DVar

	Corrective Forcing		Reynolds Forcing	
	f_{C_x}	f_{C_y}	f_{R_x}	f_{R_y}
Minimum	-1.38	-1.51	-7.85	-6.99
Maximum	1.23	0.85	3.99	9.10

The application of 3D constraints introduces a transverse forcing component f_{C_z} in addition to f_{C_x} and f_{C_y} . However, f_{C_z} remains zero across the domain and is not optimized due to the absence of mean spanwise velocity data. As shown in Tab. V, the divergence-free constraint naturally induces a spanwise velocity field that satisfies continuity, making direct control over f_{C_z} unnecessary. Including it as an independent control would lead to an under-determined system: there are five equations (continuity, three momentum equations, and the turbulence transport equation) but six unknowns (\mathbf{w} , P^* , $\tilde{\nu}$, f_{C_x} , f_{C_y} , f_{C_z}). Excluding f_{C_z} restores well-posedness under the current constraint formulation.

An important advantage of 3DVar is its ability to improve the turbulence model's accuracy while allowing the corrective forcing term to account for errors arising from the experimental setup. This is evident when comparing the magnitudes of the corrective and Reynolds forcing terms, as shown in Tab. VI, similar to the analysis conducted for 2DVar assimilation in Sec. V. While the Reynolds forcing magnitude has remained largely unchanged, the corrective forcing has significantly decreased in magnitude. This is an encouraging result, as it reinforces the idea that in 2DVar assimilation, the corrective forcing was primarily compensating for the lack of three-dimensionality in the flow due to the 2D continuity constraint. With the inclusion of a spanwise velocity component in 3DVar assimilation, the baseline RANS model is better able to capture the physics that more accurately represent

the experiment.

VII. RECONSTRUCTED QUANTITIES FOR 2DVAR AND 3DVAR USING EXPERIMENTAL REFERENCE

One key advantage of using DA is its ability to derive quantities that are not directly measured during experimental campaigns. While these quantities are not provided as reference data for assimilation, they serve as valuable benchmarks to assess the effectiveness of the assimilation method. We refer to these as reconstructed quantities or reconstructed variables and focus on them in this section. We present a comparison between the reconstructed variables from 2DVar and 3DVar assimilation, further demonstrating that 3DVar outperforms 2DVar, especially when it comes to reconstructed variables. We differentiate between physical reconstructed quantities and modeling reconstructed quantities. Among the physical reconstructed quantities, we focus on examining the pressure field and lift force, as these are of significant physical relevance and can help justify the use of assimilation. In contrast, quantities like eddy viscosity and Reynolds shear stress are reconstructed variables that hold importance within the modeling community.

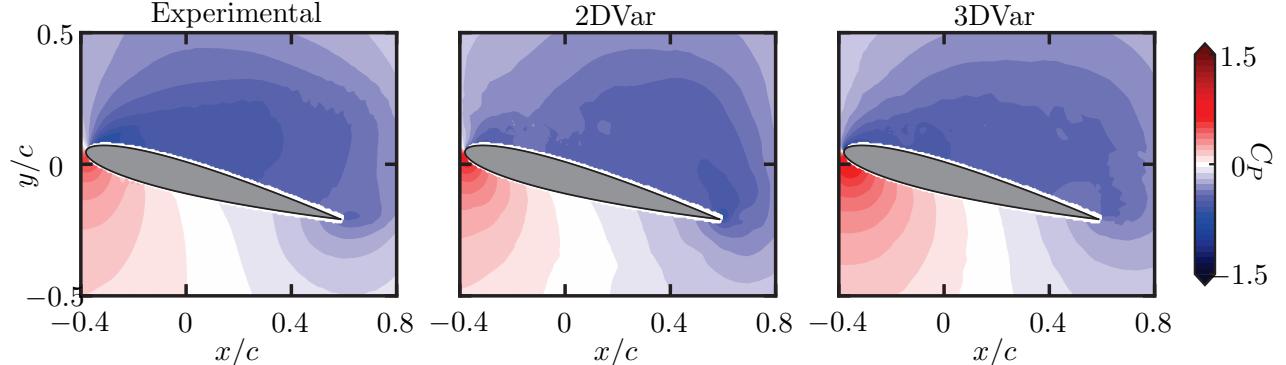


FIG. 12. Comparison of mean pressure coefficient C_P fields from experimental data, 2DVar assimilation, and 3DVar assimilation.

Pressure is a quantity that is rarely discussed in most variational DA studies due to the challenges in accurately recovering it when using the Helmholtz decomposition. We present the C_P field computed with a Poisson solver using the mean velocity and Reynolds shear stress fields from the experiment (we refer to this as Poisson pressure hereafter) and compare it with the mean pressure fields obtained from 2DVar and 3DVar assimilation in Fig. 12.

The pressure fields from 2DVar and 3DVar assimilation have been interpolated onto the experimental grid using linear interpolation. We must exercise caution when comparing the Poisson pressure and assimilated mean pressure fields. There are two key caveats to consider: first, the Poisson pressure is derived by applying a Poisson solver to the experimental mean velocity field (see Ref. [40] for details on this method). Second, the assimilated pressure field includes an irrotational component, ϕ , as dictated by the Helmholtz decomposition in Eq. 6. Therefore, direct comparison between the two pressure fields must take these differences into account.

The experimental pressure field exhibits strong suction at the leading edge on the suction side and a large positive pressure region on the pressure side. The pressure field from 2DVar does not closely resemble the experimental structure, showing a large suction region centered around the trailing edge. On the other hand, the pressure distribution obtained from 3DVar aligns more closely with the experimental field, capturing its key features more accurately. This improved reconstruction can be attributed to the enforcement of three-dimensional constraints, which allows a more physically consistent pressure field to emerge through the SIMPLE algorithm. In contrast, a majority of the corrective forcing in 2DVar is concentrated along the leading-edge shear layer (see Fig. 8), which coincides with the region of large 2D divergence error (see Fig. 3). The assimilation process compensates for this unphysical divergence using the forcing term, which in turn introduces distortions into the reconstructed pressure field. The minor differences may be attributed to the irrotational component in the Helmholtz decomposition.

In addition to examining the pressure fields, we analyze the surface pressure, a key quantity for computing the lift force. A comparison of the surface pressure coefficient between the experiment, 2DVar, and 3DVar assimilation is presented in Fig. 13. The suction peak from the experiment is not sharp, likely due to insufficient seeding at the leading edge caused by strong leading edge separation. Since a Poisson solver is used to compute the pressure field, a lack of sufficient velocity resolution affects the pressure calculation. Both 2DVar and 3DVar show a nearly constant surface pressure magnitude on the suction side, a characteristic feature of deep stall, where there is no pressure gradient along the suction side of the airfoil. The 2DVar surface pressure exhibits a small kink near the suction peak and another peak at the trailing edge. This observation aligns with the large suction in the pressure coefficient field seen in Fig. 12. The 2DVar suction-side pressure distribution does not match

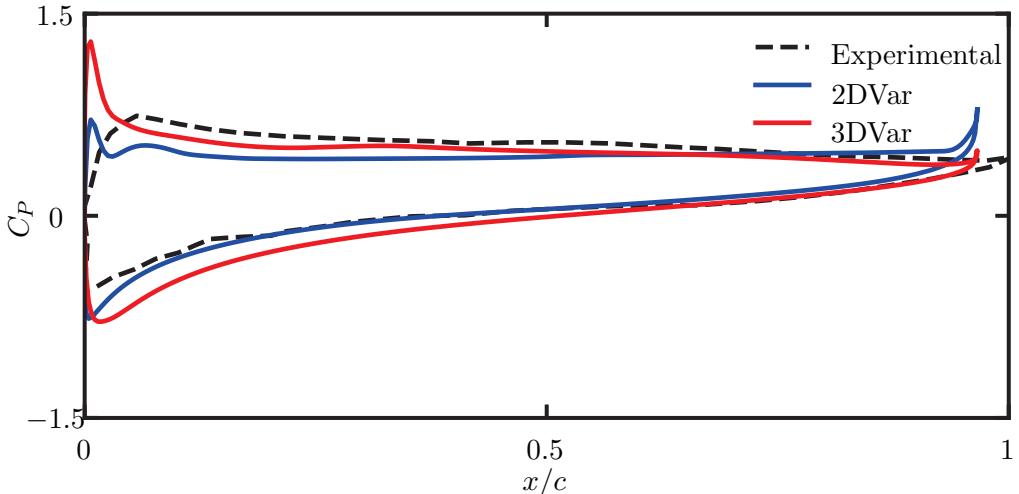


FIG. 13. Comparison of the surface mean pressure coefficients C_P from experimental data, 2DVar assimilation, and 3DVar assimilation.

TABLE VII. Comparison of lift coefficient between 2DVar and 3DVar

	Experimental	2DVar	3DVar
C_L	0.49	0.43	0.55

well with the experiment. In contrast, the 3DVar pressure distribution closely follows the experimental data, with a more pronounced suction peak.

The experimental setup includes a six-axis load cell, with further details available in Ref. [29]. The lift force measured by the load cell is reported as the lift coefficient, C_L , calculated as

$$C_L = \frac{L}{\frac{1}{2}\rho U_\infty S c}, \quad (19)$$

where S is the airfoil span. Tab. VII compares the C_L values obtained from the load cell with those from 2DVar and 3DVar assimilation. Both 2DVar and 3DVar predictions deviate by approximately 11% from the experimental value in either direction—a notable improvement over the baseline SA model ($C_L = 0.72$). However, the improvement observed with 3DVar is more consistent with the underlying flow physics, as it avoids compensating for the lack of three-dimensionality through artificial adjustments in the pressure field—a limitation that exists in 2DVar assimilation. This highlights the advantage of assimilation, as it enables the

estimation of lift force without requiring additional measurement devices.

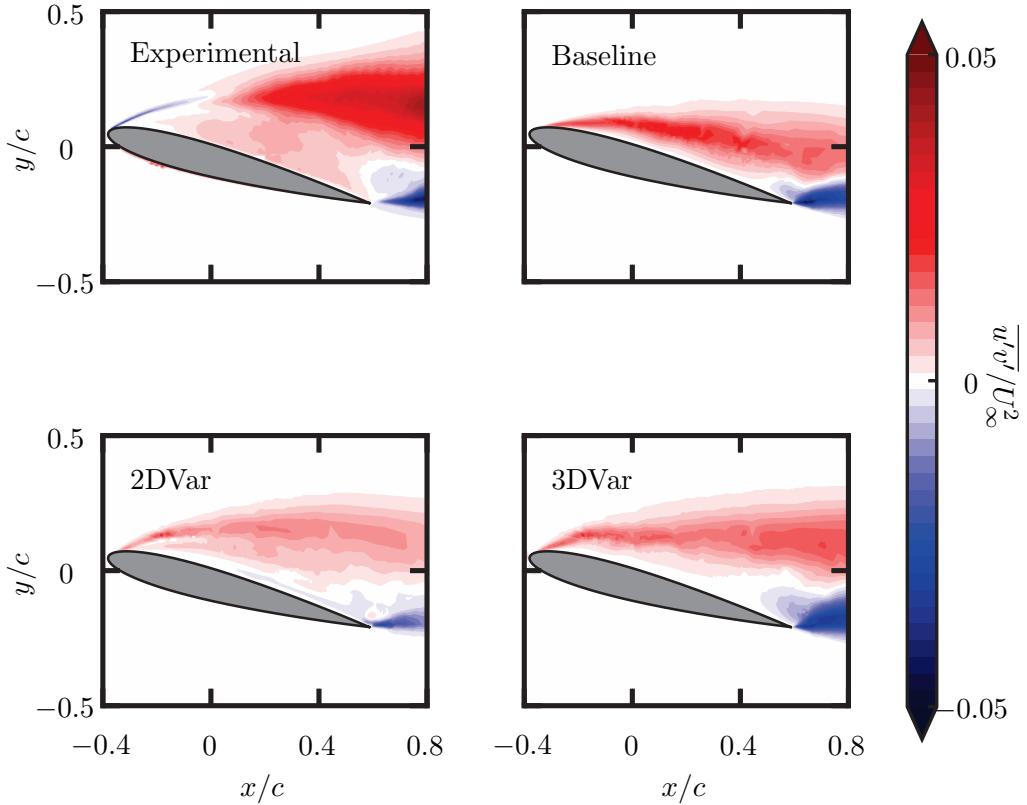


FIG. 14. Comparison of Reynolds shear stress $\overline{u'v'}/U_\infty^2$ across experimental reference data, baseline SA RANS model, 2DVar, and 3DVar assimilation.

We have examined the benefits of 3DVar over 2DVar in reconstructing physically relevant quantities. Now, we shift our focus to reconstructed quantities that hold significance from a modeling perspective. Specifically, we compare the Reynolds shear stress for 2DVar and 3DVar assimilation against the experimental reference in Fig. 14. It is important to note that the Reynolds shear stress shown here is derived solely from the turbulence model, excluding the contribution from the corrective forcing. This choice is intentional. As mentioned before, one of our goals is to reassess the role of the forcing term—not simply as a universal corrector, but as a means to absorb data inconsistencies and experimental imperfections, especially in regions where the flow exhibits significant three-dimensionality.

In the case of 3DVar, the presence of spanwise velocity gradients and 3D constraints allows the turbulence model to account for the physics of deep stall more directly, while the corrective forcing plays a secondary role, primarily addressing experimental uncertainties. This is supported by the reduced magnitude of the corrective forcing in Tab. VI. In contrast,

for 2DVar, the absence of 3D constraints forces the corrective term to compensate not only for modeling limitations but also for unmodeled physical effects, such as out-of-plane transport. This results in higher forcing magnitudes and potentially unphysical reconstructions of quantities like pressure and eddy viscosity. By presenting the Reynolds shear stress in this form, we aim to emphasize how enabling physical three-dimensionality shifts more of the reconstruction burden onto the turbulence model itself, yielding a more meaningful and interpretable correction. For completeness, we also include the Reynolds shear stress field from the baseline model for comparison.

The baseline model exhibits a positive Reynolds shear stress region along the leading-edge shear layer, which is displaced downward and remains closer to the airfoil surface relative to the experimental reference. This suggests a smaller recirculation bubble, as corroborated by the $U_x = 0$ contour in Fig. 5. In contrast, the Reynolds shear stress distributions obtained from both 2DVar and 3DVar assimilation align more closely with the experiment, where the positive Reynolds shear stress region near the leading edge is displaced upward. While both assimilation approaches show improvement, we further investigate the differences observed between 2DVar and 3DVar assimilation.

We expand the field of view in the streamwise direction to highlight better the differences in Reynolds shear stress between 2DVar and 3DVar assimilation. Since we employ an SA RANS baseline model, we also analyze the eddy viscosity ν_t , which the SA model utilizes to relate the Reynolds stresses to the mean rate of strain obtained from velocity gradients. The scaled eddy viscosity, ν_t/ν , and the Reynolds shear stress, R_{12}/U_∞^2 , scaled by the freestream velocity, are presented in Fig. 15. The majority of the eddy viscosity for 2DVar remains concentrated within the recirculation bubble. In contrast, the eddy viscosity field for 3DVar extends further into the wake of the airfoil. Another key distinction is that the eddy viscosity is significantly stronger for 3DVar than for 2DVar. A stronger eddy viscosity enhances mixing and implies a greater transport of momentum. This trend is also reflected in the distribution of Reynolds shear stress in Fig. 15, where the positive component near the leading edge shifts further into the wake, while the negative component at the trailing edge becomes more pronounced. Additionally, the Reynolds shear stress extends further downstream for 3DVar than for 2DVar, indicating a more sustained turbulence effect in the wake and suggesting enhanced momentum transfer over a larger spatial extent. These observations make it clear that the Reynolds shear stress from 3DVar shows a better agreement with the experiment

than 2DVar.

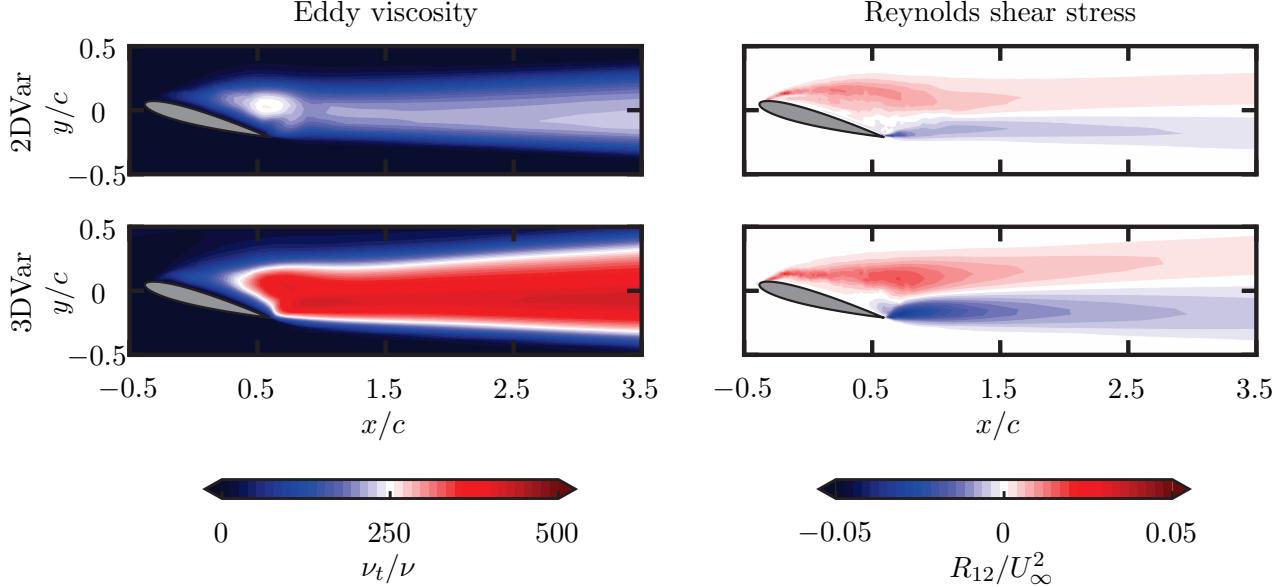


FIG. 15. Scaled eddy viscosity ν_t/ν and Reynolds shear stress comparison between 2DVar and 3DVar assimilation using experimental reference data.

This improvement in the agreement of the Reynolds shear stress from 3DVar and the enhancement of the eddy viscosity in the wake of the airfoil highlight the ability of the baseline turbulence model formulation to capture complex flow physics, such as deep stall at high Reynolds numbers. It provides clear evidence that assimilating inherently 3D flows with 2D constraints distorts the balance of contributions, shifting more of the burden away from the existing turbulence model. Importantly, we have not altered the functional form of the turbulence model. As previously mentioned, the inclusion of a corrective forcing term in the momentum equations, along with 3D constraints, allows the turbulence model to function properly, while the forcing term compensates for deficiencies in the experimental setup. This has significant implications for both the physically relevant quantities (such as surface pressure and lift forces) and those that are critical from a modeling perspective (such as eddy viscosity and Reynolds stress tensor).

VIII. CONCLUSION

We present a novel approach for assimilating 2D2C time-averaged velocity fields obtained from experimental measurements using variational DA with a discrete adjoint method. By

enforcing constraints in three dimensions (3DVar), we demonstrate its benefits over conventional variational DA with 2D constraints (2DVar). In both cases, the control variable is introduced as a corrective forcing term in the momentum equations. We apply this method to a deep stall case of a NACA0012 airfoil at a moderate Reynolds number, characterized by strong leading-edge flow separation. Our results show that assimilating an experimental velocity field that is not divergence-free using 2D constraints leads to an increased magnitude of the corrective forcing term, as it compensates for the lack of divergence in the data. This distortion disrupts the balance of the momentum equations and weakens the role of the turbulence model. The true contribution of the corrective forcing term is difficult to isolate, as it accounts for both discrepancies arising from the RANS turbulence model formulation and inconsistencies in the experimental data, including measurement errors, noise, and incomplete information.

The ambiguity in interpreting the corrective forcing is resolved by implementing 3DVar, where the flow is allowed to develop in the spanwise direction. This enables a more physically consistent representation of the 3D nature of the flow field while correctly enforcing 3D constraints. We demonstrate that the divergence error in the experimental velocity field is compensated by the spanwise velocity gradient, effectively accounting for missing data in the experiment that caused the divergence error. The resulting corrective forcing field primarily addresses artifacts in the experimental setup, shifting the balance toward the turbulence model through an enhanced Reynolds forcing. Although 3DVar may not always result in lower data mismatch compared to 2DVar, it enforces three-dimensional constraints that are physically consistent with the nature of a separated, three-dimensional flow. This ensures that the reconstructed unmeasured quantities, such as pressure and eddy viscosity, do not absorb deficiencies arising from the limitations of 2D constraints, leading to more meaningful and physically interpretable results.

While our 3DVar implementation utilized 2D2C experimental reference data, one of the key outcomes of this study was demonstrating the strengths of 3DVar in assimilating complex fluid dynamics problems at flow conditions of practical relevance. As future work, we plan to extend this approach by incorporating three-component mean velocity fields for assimilation using experimental data. There remains limited clarity on the data requirements for such problems, and few studies have examined how reconstructed quantities are influenced by the amount of data used. Additionally, we aim to perform uncertainty quantification on

the obtained corrective forcing, as this remains a topic of significant interest within the DA community.

ACKNOWLEDGMENTS

The authors acknowledge the use of the IRIDIS High-Performance Computing Facility and its support services at the University of Southampton. We gratefully acknowledge funding from EPSRC (Grant Ref: EP/W009935/1) and the School of Engineering at the University of Southampton for UCP’s PhD studentship. The authors declare no conflicts of interest. UCP contributed to the development of DA algorithms, conducted computations, analyzed the data, and drafted multiple versions of the manuscript. BG and SS were responsible for conceptualization, funding acquisition, manuscript editing, and project management.

Appendix A: Projection and smoothing

The details of the implementation of projection and smoothing operations are presented here. We consider a field measurement technique such as planar PIV where the output of the processing step is a mean-velocity field defined on a regular cartesian grid of uniform spacing. This is illustrated in Fig. 16 as blue crosses. Let us assume that the experimental data co-ordinates are cell centroids of some fictitious cell of uniform dimensions shown as dashed lines. The computational mesh that is used by the DA algorithm is represented as an irregular grid with non-uniform cell size and is superimposed on the regular grid. The 2D representation of this mesh is shown here since each cell is actually a hexahedral element.

Consider one such cartesian grid cell that is outlined in red and enlarged for the sake of explanation. This cell encloses the cell-centroids of four computational mesh cells that have been highlighted in grey. The goal of the algorithm is to find the volume-averaged velocity at the centroid of the cartesian grid cell using the contribution of all the computational mesh cells that lie inside it. Mathematically, this is given by:

$$\bar{u}_g = \frac{\sum_{i=1}^N u_i V_i}{\sum_{i=1}^N V_i}, \quad (A1)$$

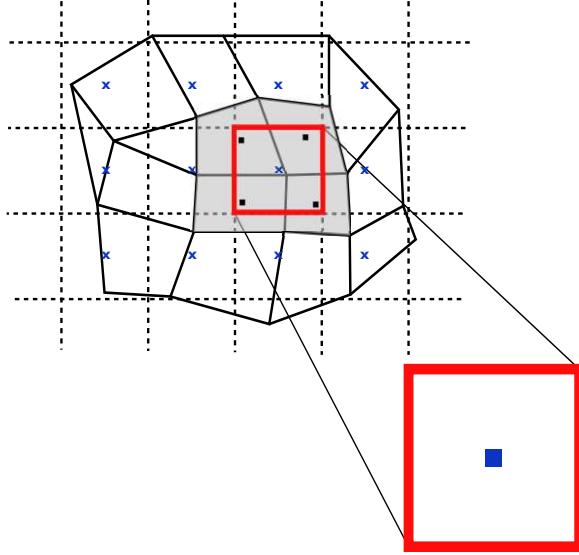


FIG. 16. Graphical illustration of the cell volume weighted averaging. A structured grid (---) with experimental data points (x) superimposed on an unstructured grid (—) is shown with cell centroids of the CFD mesh (■) and cell volume weighted grid (■).

where \bar{u}_g is the volume-averaged velocity at the grid cell centroid, N is the number of computational mesh cell centroids that lie within a given grid cell (in this example, 4), u_i and V_i are the streamwise velocity and volume of the i^{th} mesh cell, respectively.

The cell-volume weighted averaging procedure can be thought of as the projection operator \mathcal{Q} defined in Eq. 7. The output of this process is the projected computational variables used for calculating the objective function. The objective function is calculated on the regular Cartesian grid. As mentioned before, the adjoint solver operates on the computational mesh used in the RANS simulation. This requires a second projection $\hat{\mathcal{Q}}(.)$ (which we will call smoothing) of the objective function onto the computational mesh. The second projection (or smoothing operator) serves as the discrete analogue of the adjoint of the projection operator. This step can also be interpreted as a way of computing the pseudo-inverse of the projection operator. As suggested in [13], the objective function at a given Cartesian grid point is distributed equally among all the computational cells that lie inside it.

There is a small modification when performing the projection and smoothing operations for 3D DA. This entails a search in the spanwise direction within the bounds set by the user. Typically, the search is limited to a single layer of cells along the spanwise direction by specifying z_{min} and z_{max} where $\Delta z = z_{max} - z_{min} = /N_z$, where N_z is the number of points

in the spanwise direction. The choice of N_z should then be approximately representative of the laser sheet thickness in a PIV setup. The volume averaging is then restricted to this single layer to avoid the ambiguity of projecting and smoothing more than a single layer of cells.

- [1] F.-X. Le Dimet and O. Talagrand, Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects, *Tellus A: Dyn. Meteorol. Oceanogr.* **38**, 97 (1986).
- [2] D. P. G. Foures, N. Dovetta, D. Sipp, and P. J. Schmid, A data-assimilation method for Reynolds-Averaged Navier–Stokes-driven mean flow reconstruction, *J. Fluid Mech.* **759**, 404 (2014).
- [3] S. Symon, D. Sipp, and B. J. McKeon, A tale of two airfoils: resolvent-based modelling of an oscillator versus an amplifier from an experimental mean, *J. Fluid Mech.* **881**, 51 (2019).
- [4] L. Franceschini, D. Sipp, and O. Marquet, Mean-flow data assimilation based on minimal correction of turbulence models: Application to turbulent high Reynolds number backward-facing step, *Phys. Rev. Fluids* **5**, 094603 (2020).
- [5] C. Thompson, U. Cadambi Padmanaban, B. Ganapathisubramani, and S. Symon, The effect of variations in experimental and computational fidelity on data assimilation approaches, *Theor. Comput. Fluid. Dyn.* **38**, 431 (2024).
- [6] V. Mons, A. Vervynck, and O. Marquet, Data assimilation and linear analysis with turbulence modelling: application to airfoil stall flows with PIV measurements, *Theor. Comput. Fluid. Dyn.* **38**, 403 (2024).
- [7] S. B. Pope, Turbulent flows, *Meas. Sci. Technol.* **12**, 2020 (2001).
- [8] R. J. Adrian, Particle-imaging techniques for experimental fluid mechanics, *Annu. Rev. Fluid. Mech.* **23**, 261 (1991).
- [9] D. Barkley and R. D. Henderson, Three-dimensional Floquet stability analysis of the wake of a circular cylinder, *J. Fluid Mech.* **322**, 215 (1996).
- [10] S. Gupta, J. Zhao, A. Sharma, A. Agrawal, K. Hourigan, and M. C. Thompson, Two-and three-dimensional wake transitions of a NACA0012 airfoil, *J. Fluid Mech.* **954**, A26 (2023).
- [11] J. E. Peter and R. P. Dwight, Numerical sensitivity analysis for aerodynamic optimization: A

survey of approaches, *Comput. Fluids* **39**, 373 (2010).

[12] G. K. Kenway, C. A. Mader, P. He, and J. R. Martins, Effective adjoint approaches for computational fluid dynamics, *Prog. Aerosp. Sci.* **110**, 100542 (2019).

[13] S. Symon, N. Dovetta, B. J. McKeon, D. Sipp, and P. J. Schmid, Data assimilation of mean velocity from 2D PIV measurements of flow over an idealized airfoil, *Exp. Fluids* **58**, 1 (2017).

[14] O. Brenner, P. Piroozmand, and P. Jenny, Efficient assimilation of sparse data into RANS-based turbulent flow simulations using a discrete adjoint method, *J. Comput. Phys.* **471**, 111667 (2022).

[15] O. Brenner, J. Plogmann, P. Piroozmand, and P. Jenny, A variational data assimilation approach for sparse velocity reference data in coarse RANS simulations through a corrective forcing term, *Comput. Methods Appl. Mech. Eng.* **427**, 117026 (2024).

[16] A. S. Cato, P. S. Volpiani, V. Mons, O. Marquet, and D. Sipp, Comparison of different data-assimilation approaches to augment RANS turbulence models, *Comput. Fluids* , 106054 (2023).

[17] P. He, C. A. Mader, J. R. Martins, and K. J. Maki, DAFoam: An open-source adjoint framework for multidisciplinary design optimization with openfoam, *AIAA J.* **58**, 1304 (2020).

[18] P. He, C. A. Mader, J. R. Martins, and K. J. Maki, An aerodynamic design optimization framework using a discrete adjoint approach with OpenFOAM, *Comput. Fluids* **168**, 285 (2018).

[19] H. G. Weller, G. Tabor, H. Jasak, and C. Fureby, A tensorial approach to computational continuum mechanics using object-oriented techniques, *Comput. Phys.* **12**, 620 (1998).

[20] U. Cadambi Padmanaban, B. Ganapathisubramani, C. Vanderwel, and S. Symon, Towards passive scalar reconstruction using data assimilation, in *Proceedings of the 14th UK Conference on Wind Engineering: University of Southampton* (2024).

[21] M. Raissi, P. Perdikaris, and G. E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J. Comput. Phys.* **378**, 686 (2019).

[22] L. Sliwinski and G. Rigas, Mean flow reconstruction of unsteady flows using physics-informed neural networks, *Data-Centric Eng.* **4**, e4 (2023).

[23] Y. Patel, V. Mons, O. Marquet, and G. Rigas, Turbulence model augmented physics-informed neural networks for mean-flow reconstruction, *Phys. Rev. Fluids* **9**, 034605 (2024).

[24] J. G. von Saldern, J. M. Reumschüssel, T. L. Kaiser, M. Sieber, and K. Oberleithner, Mean flow data assimilation based on physics-informed neural networks, *Phys. Fluids* **34** (2022).

[25] A. Villié, S. Schmitter, J. G. von Saldern, S. Demange, and K. Oberleithner, Physics-informed neural networks for enhancing medical flow magnetic resonance imaging: Artifact correction and mean pressure and reynolds stresses assimilation, *Phys. Fluids* **37** (2025).

[26] C. Toma, B. Ganapathisubramani, and S. Symon, Mixed data-source transfer-learning for a turbulence model augmented physics-informed neural network., under review (2025).

[27] V. Mons, J.-C. Chassaing, T. Gomez, and P. Sagaut, Reconstruction of unsteady viscous flows using data assimilation schemes, *J. Comput. Phys.* **316**, 255 (2016).

[28] P. S. Volpiani, M. Meyer, L. Franceschini, J. Dandois, F. Renac, E. Martin, O. Marquet, and D. Sipp, Machine learning-augmented turbulence modeling for RANS simulations of massively separated flows, *Phys. Rev. Fluids* **6**, 064607 (2021).

[29] D. W. Carter and B. Ganapathisubramani, Low-order modeling and sensor-based prediction of stalled airfoils at moderate reynolds number, *AIAA J.* **61**, 2893 (2023).

[30] P. Spalart and S. Allmaras, A one-equation turbulence model for aerodynamic flows, in *30th aerospace sciences meeting and exhibit* (1992) p. 439.

[31] K. C. Giannakoglou and D. I. Papadimitriou, Adjoint methods for shape optimization, *Optim. Comput. Fluid Dyn.* , 79 (2008).

[32] N. G. M. Sagebaum, T. Albring, High-performance derivative computations using codipack, *ACM Trans. Math. Softw.* **45** (2019).

[33] R. B. Wilson, A simplicial algorithm for concave programming, Ph. D. Dissertation, Graduate School of Bussiness Administration (1963).

[34] P. E. Gill, W. Murray, and M. A. Saunders, SNOPT: An SQP algorithm for large-scale constrained optimization, *SIAM review* **47**, 99 (2005).

[35] M. Raffel, C. E. Willert, F. Scarano, C. J. Kähler, S. T. Wereley, and J. Kompenhans, *Particle image velocimetry: a practical guide* (Springer, 2018).

[36] A. Sciacchitano, Uncertainty quantification in particle image velocimetry, *Meas. Sci. Technol.* **30**, 092001 (2019).

[37] R. F. Huang and C. L. Lin, Vortex shedding and shear-layer instability of wing at low-reynolds numbers, *AIAA J.* **33**, 1398 (1995).

[38] S. Yarusevych and M. S. Boutilier, Vortex shedding of an airfoil at low reynolds numbers,

AIAA. J **49**, 2221 (2011).

[39] S. Patankar and D. Spalding, A calculation procedure for heat, mass and momentum transfer in three-dimensional parabolic flows, *Int. J. Heat Mass Transf.* **15**, 1787 (1972).

[40] R. De Kat and B. Van Oudheusden, Instantaneous planar pressure determination from PIV in turbulent flow, *Exp. Fluids* **52**, 1089 (2012).

[41] O. Marquet, V. Mons, M. Zauner, and B. Leclaire, Turbulent mean flow estimation with state observer assimilation of velocity measurements in RANS models, in *TSFP 12* (2022).

[42] C. He, X. Zeng, P. Wang, X. Wen, and Y. Liu, Four-dimensional variational data assimilation of a turbulent jet for super-temporal-resolution reconstruction, *J. Fluid Mech.* **978**, A14 (2024).

[43] A. H. Gebremedhin, F. Manne, and A. Pothen, What color is your Jacobian? Graph coloring for computing derivatives, *SIAM review* **47**, 629 (2005).

[44] W. P. Jones and B. E. Launder, The prediction of laminarization with a two-equation model of turbulence, *Int. J. Heat Mass Transf.* **15**, 301 (1972).

[45] D. C. Wilcox, Reassessment of the scale-determining equation for advanced turbulence models, *AIAA J.* **26**, 1299 (1988).

[46] C. M. de Silva, J. Philip, and I. Marusic, Minimization of divergence error in volumetric velocity measurements and implications for turbulence statistics, *Exp. Fluids* **54**, 1 (2013).

[47] T. Hayase, Numerical simulation of real-world flows, *Fluid Dyn. Res.* **47**, 051201 (2015).

[48] F. R. Menter, Two-equation eddy-viscosity turbulence models for engineering applications, *AIAA J.* **32**, 1598 (1994).

[49] D. J. S. Aulakh, X. Yang, and R. Maulik, Robust experimental data assimilation for the Spalart-Allmaras turbulence model, *Phys. Rev. Fluids* **9**, 084608 (2024).

[50] M. Zauner, V. Mons, O. Marquet, and B. Leclaire, Nudging-based data assimilation of the turbulent flow around a square cylinder, *J. Fluid Mech.* **937**, A38 (2022).