

CONTINUOUS UNSTEADY ADJOINT VARIATIONAL MULTISCALE METHOD FOR AIRFOILS SHAPE OPTIMIZATION

Carlo Brunelli[1,2,3], Matija Avirovic[1,3,4], Bart Janssens[1] & Benoît G. Marinus[1]

[1]Royal Military Academy, Brussels, Belgium
[2]Vrije Universiteit Brussel, Brussels, Belgium
[3]von Karman Institute for Fluid Dynamics, Sint-Genesius-Rode, Belgium
[4]Ghent University, Gent, Belgium

Abstract

Aerodynamic shape optimization in unsteady viscous flows remains challenging, as the adjoint equations must be solved backward in time and require stabilization consistent with the forward flow solver. Most existing continuous adjoint formulations rely on SUPG/PSPG stabilization, which, while effective, lacks the theoretical generality needed to systematically extend the method to more complex flow regimes. In this work, a novel shape optimization framework is presented in which the continuous unsteady adjoint equations are derived within the Variational Multiscale (VMS) framework, providing a more consistent and extensible foundation for both the flow and adjoint solvers. Mesh deformation is handled via a Radial Basis Function (RBF) scheme, which preserves local geometric features throughout the optimization cycle, and the L-BFGS algorithm is employed to drive the minimization of the objective function.

The method is validated and tested on a two-dimensional NACA 0012 airfoil at Reynolds number $Re = 1000$. Gradient accuracy is first assessed at an angle of attack of $2.5°$, showing good agreement with finite difference estimates. At an angle of attack of $4°$, inverse design cases targeting time-averaged lift coefficients of $\bar{C}_L = 0.4$, 0.5 and 0.6 are conducted. In all cases, the optimizer successfully converges toward the prescribed target, with the gradient norm decreasing by three orders of magnitude, confirming approach toward a local minimum. These results demonstrate the capability and reliability of the proposed VMS-based adjoint framework for aerodynamic shape optimization in unsteady flow conditions.

Keywords: Adjoint, Variational MultiScale Method, Airfoil, Optimization

1. Introduction

1.1 Background and Motivation

The adjoint method is a well-established and efficient methodology for evaluating sensitivities for CFD-based aerodynamic design. Airfoil shape optimization using the adjoint method to compute sensitivities with respect to shape design parameters has been pursued for more than three decades [1, 2, 3, 4, 5, 6, 7, 8, 9, 10]. A comprehensive treatment of adjoint methods in the context of shape optimization is provided in [11]. Unsteady adjoint optimization is challenging but provides accurate gradients for flow problems where traditional steady-state assumptions break down [12, 13]. For time-averaged objective functions, the adjoint equations must be integrated backward in time over the entire primal simulation window, which imposes significant memory and computational requirements. Adjoint solvers are not always available for traditional RANS models such as the $k-\omega$ SST with $\gamma - Re_\theta$ transition modeling. The constitutive empirical correlations of these models are often non-differentiable, requiring ad-hoc regularization strategies to enable adjoint derivation [14], and do not permit a fully consistent formulation of the adjoint system. However, some regularization of the correlations have been carried out [15], leading to smooth approximations of the non-differentiable source terms and empirical correlations without degrading the accuracy of the transition predictions.

This demonstrates that consistent, differentiable transition modeling is achievable, and the resulting smooth formulation is well suited to gradient-based optimization frameworks that rely on adjoint sensitivities. This limitation is most critical in regimes characterized by Laminar Separation Bubbles (LSBs), separation, or laminar–turbulent transition, where RANS closures typically fail to provide reliable results [16, 17], and where a physically consistent adjoint formulation is precisely most needed.

1.2 Present Contribution

This research employs the Variational Multiscale (VMS) framework [18, 19] for Large Eddy Simulation (LES) to solve both the primal and adjoint flow problems [20, 21]. The main novelty of this study is the derivation and application of the adjoint equations directly from the VMS formulation, making the approach suitable for unsteady flow regimes and laminar–turbulent transition modeling. An adjoint formulation is incorporated within the LES framework to optimize time-averaged objective functions, constituting a novel methodological advancement. The continuous adjoint of the stabilized finite element Navier–Stokes formulation was first explored in [22]; the present work extends this contribution to the full VMS-LES setting, incorporating the effect of the subscale model in the sensitivity computation and enabling the treatment of genuinely unsteady, transitional flows. A central feature of the proposed approach is that the adjoint equations are derived directly and consistently from the same VMS formulation used for the primal problem. As a result, the primal and adjoint solutions share the same discretization framework, the same stabilization operators, and the same numerical scheme. This internal consistency is essential for the reliability of the computed sensitivities, as any mismatch between the numerics of the primal and adjoint systems would introduce spurious gradient errors, potentially corrupting the optimization trajectory. The feasibility of the approach at the moderate Reynolds numbers and two-dimensional configurations considered here makes this cost tractable, while the accuracy of the resulting sensitivities is fully retained. The outcome is an optimization framework that fully respects the governing equations at all scales resolved by the LES. In these conditions, the VMS-LES solver provides a physically accurate primal solution, while the consistent derivation of the adjoint equations within the same framework guarantees that the computed sensitivities faithfully represent the response of the unsteady flow to geometric perturbations. This is precisely the regime, involving transition and separation, where the present approach offers the most significant advantage over RANS-based methods. The methodology proposed here enables shape optimization in regimes where RANS-based adjoint approaches are either unavailable or lack robustness. Radial Basis Functions (RBF) are used to morph both the airfoil geometry and the computational mesh [23, 24]. The underlying VMS-LES solver, described in detail in [25, 26], has been validated for low-Reynolds airfoil configurations in [16, 17, 27, 28], providing confidence in the fidelity of the primal flow solutions upon which the adjoint is built.

1.3 Outline

The research paper is organized as follows. The primal VMS-LES formulation is briefly recalled, followed by the derivation of the continuous unsteady adjoint equations within the same framework and using the same numerical scheme. The adjoint solver is then validated against finite difference gradient estimates on a NACA0012 airfoil at $Re = 1000$ and angle of attack $\alpha = 2.5^\circ$, demonstrating excellent agreement and confirming the correctness of the sensitivity computation. Finally, the full optimization loop is applied and validated, successfully targeting a time-averaged lift coefficient $\overline{C_L} = 0.60$. Together, these results open a new stage in adjoint-based aerodynamic optimization, showing that consistent, high-fidelity unsteady adjoint methods are now feasible within a unified VMS-LES framework, combining physical accuracy in the primal solution with rigorous, numerically coherent sensitivity analysis.

2. VMS-Consistent Continuous Adjoint Formulation

2.1 Adjoint framework and objective gradient

The continuous adjoint method is employed to compute efficiently the sensitivity of a cost functional I_C, or objective function, with respect to a set of geometric parameters $\boldsymbol{\beta}$. The objective depends also on the incompressible flow solution $I_C(\boldsymbol{U}, \beta)$, where $\mathbf{U} = (\boldsymbol{u}, p)$ and typically represents integrated

aerodynamic quantities such as drag, lift, or efficiency. An augmented functional is introduced to enforce the governing equations through Lagrange multipliers,

$$I = I_C + \int_{t_0}^{t_F} \int_{\Omega} \mathbf{\Psi} \cdot \mathscr{R}(\mathbf{U}) \, \mathrm{d}\Omega \, \mathrm{d}t, \tag{1}$$

where $\mathscr{R}$ denotes the incompressible Navier–Stokes operator and $\mathbf{\Psi} = (\psi_u, \psi_p)$ are the adjoint velocity and pressure. When the primal equations are satisfied, I is equivalent to I_C.
Differentiation of I with respect to β yields the total gradient:

$$\frac{dI}{d\beta} = \frac{\partial I}{\partial \beta} + \frac{\partial I}{\partial \mathbf{U}} \frac{\partial \mathbf{U}}{\partial \beta} + \frac{\partial I}{\partial \mathbf{\Psi}} \frac{\partial \mathbf{\Psi}}{\partial \beta} \tag{2}$$

where the dependence on $\frac{\partial I}{\partial \mathbf{\Psi}}$ and $\frac{\partial I}{\partial \mathbf{U}}$ are eliminated by enforcing the primal (4) and adjoint (3) equations:

$$\frac{\partial I}{\partial \mathbf{U}} = \frac{\partial I_C}{\partial \mathbf{U}} + \frac{1}{\Delta T} \frac{\partial}{\partial \mathbf{U}} \int_{t_0}^{t_F} \int_{\Omega} \mathbf{\Psi} \cdot \mathscr{R}(\mathbf{U}) d\Omega dt = 0 \tag{3}$$

$$\frac{\partial I}{\partial \mathbf{\Psi}} = \frac{\partial I_C}{\partial \mathbf{\Psi}} + \frac{1}{\Delta T} \frac{\partial}{\partial \mathbf{\Psi}} \int_{t_0}^{t_F} \int_{\Omega} \mathbf{\Psi} \cdot \mathscr{R}(\mathbf{U}) d\Omega dt = \int_{t_0}^{t_F} \int_{\Omega} \mathscr{R}(\mathbf{U}) d\Omega dt = 0 \tag{4}$$

As a result, the gradient depends only on explicit geometric terms and boundary contributions, enabling cost-independent sensitivity evaluation.
The key contribution of this work is a VMS-consistent adjoint formulation, obtained by applying the adjoint operator to the same variational multiscale decomposition used to discretize the primal flow. This ensures operator consistency and preserves LES-level sensitivity fidelity.

2.2 VMS formulation of the primal problem

The unsteady incompressible Navier–Stokes equations are discretized using a Variational Multiscale (VMS) finite element method. The solution is decomposed into resolved and unresolved components, with modeling confined to the subgrid scales through residual-based terms. This approach provides implicit turbulence modeling without empirical parameters and avoids the need for wall functions.
Let $\mathbf{V}_h \times Q_h$ denote equal-order conforming finite element spaces for velocity and pressure. The Galerkin weak form is augmented by extra terms:

$$B_G(t; \mathbf{u}_h, p_h; \mathbf{v}_h, q_h) + B_{STAB}(t; \mathbf{u}_h, p_h; \mathbf{v}_h, q_h) + B_{VMS}(t; \mathbf{u}_h, p_h; \mathbf{v}_h, q_h) = 0 \qquad \forall (\mathbf{v}_h, q_h) \in \mathbf{V}_h \times Q_h, \tag{5}$$

where B_{STAB} includes residual-based SUPG, PSPG, and LSIC operators, and B_{VMS} accounts for cross-scale interactions driven by the momentum and continuity residuals. Stabilization parameters depend on the local mesh metric and time step, following standard VMS practice. See [21] and [26] for more details on $B_{\mathrm{STAB}}, B_{\mathrm{VMS}}$ terms and how they are implemented. The results is a set of time-dependent PDEs.

2.3 Continuous time dependent VMS adjoint

The adjoint problem is derived by taking the first variation of the augmented functional with respect to the primal variables. Using adjoint test spaces consistent with the primal discretization $(\mathbf{v}_{\dagger_h}, q_{\dagger_h})$, a linear time-dependent adjoint system is obtained. The adjoint momentum and continuity residuals are:

$$\mathbf{R}_m^{\dagger} = -\partial_t \phi_u - \nabla \phi_p - (\nabla \phi_u)^T \mathbf{u}_h - (\nabla \phi_u) \mathbf{u}_h - \nu \Delta \phi_u, \tag{6}$$

$$R_c^{\dagger} = -\nabla \cdot \phi_u. \tag{7}$$

The adjoint variational formulation mirrors the primal one and includes the adjoint counterparts of both stabilization and VMS terms (8), (9). This yields a consistent dual operator and avoids ad hoc adjoint modeling. The adjoint system is solved backward in time and requires storage of the primal trajectory.

$$B_{G_\dagger} = \int_\Omega \left(\nu \cdot \nabla(\mathbf{v}_{\dagger_h}) \odot \nabla(\phi_{\mathbf{u}_h}) - q_{\dagger_h}(\nabla \cdot \phi_{\mathbf{u}_h}) \right) d\Omega + \int_\Omega \left(\mathbf{v}_{\dagger_h} \odot \left(\frac{\phi_{\mathbf{u}_h}}{\partial t} - \nabla(\phi_{p_h}) - \nabla(\phi_{\mathbf{u}_h})^T \cdot \mathbf{u}_h - \nabla(\phi_{\mathbf{u}_h}) \cdot \mathbf{u}_h \right) \right) d\Omega \tag{8}$$

$$B_{STAB_\dagger} = \int_\Omega \left((-{\mathbf{v}_{\dagger_h}}^T \cdot \mathbf{u}_h - \mathbf{v}_{\dagger_h} \cdot \mathbf{u}_h - \nabla(q_{\dagger_h})) \odot \tau_m \mathbf{R}_{\mathbf{madj}} \right) d\Omega - \int_\Omega \left(\tau_c \cdot (\nabla \cdot \mathbf{v}_{\dagger_h}) \odot R_{cadj} \right) d\Omega \tag{9}$$

The new adjoint terms, arising from the VMS are:

$$B_{VMS1_\dagger} = \int_\Omega \left((-\mathbf{u}_h \cdot ({\mathbf{v}_{\dagger_h}}^T) \odot \tau_m \mathbf{R}_{\mathbf{madj}} \right) d\Omega \tag{10}$$

Boundary conditions for the adjoint variables follow from the objective functional and duality arguments. On solid boundaries, they depend explicitly on the aerodynamic quantity of interest, while homogeneous conditions are imposed on far-field boundaries to ensure well-posedness and numerical stability.

This VMS-consistent continuous adjoint enables robust and accurate gradient evaluation for unsteady turbulent flows, forming the foundation for high-fidelity shape optimization at practical Reynolds numbers.

3. Radial Basis Function Geometry Deformation

Airfoil shape changes are performed using a Radial Basis Function (RBF) mesh deformation technique, which is well suited for adjoint-based shape optimization. RBF-based morphing was originally introduced for aerodynamic mesh deformation in [23] and has since been successfully applied to airfoil and wing optimization problems, demonstrating robustness and high mesh quality [24]. Geometry modifications are prescribed through a limited set of surface control points whose displacements are determined by the optimization algorithm, driven by gradients obtained from the adjoint solution. The resulting deformation is smoothly propagated to the surrounding mesh, preserving mesh quality without requiring re-meshing.

Starting from an initial grid configuration $\{\mathbf{x}_i^0\}$, the deformed grid is obtained as

$$\mathbf{x}_i = \mathbf{x}_i^0 + \Delta\mathbf{x}_i,$$

where the displacement field $\Delta\mathbf{x}_i$ is reconstructed using RBF interpolation. Boundary nodes are divided into moving nodes, where displacements are prescribed, and fixed nodes, where zero displacement is enforced to anchor the deformation, following the classical formulation of [23].

The continuous displacement field is defined as

$$\mathbf{s}(\mathbf{x}) = \sum_{j=1}^{N_b} \boldsymbol{\alpha}_j \, \phi(\|\mathbf{x} - \mathbf{x}_j\|) + \mathbf{P}(\mathbf{x}),$$

with ϕ a radial basis function, $\boldsymbol{\alpha}_j$ the interpolation coefficients, and $\mathbf{P}$ a low-order polynomial ensuring rigid-body consistency. The coefficients are obtained by solving a small linear system enforcing the prescribed boundary displacements, as commonly done in RBF mesh-morphing techniques [23, 24]. Compactly supported RBFs are employed to balance smoothness and computational efficiency, limiting the influence of each control point to a finite radius. This yields sparse interpolation matrices while still providing smooth and physically meaningful deformations, an approach that has been shown to generate high-quality aerodynamic meshes for airfoil configurations [24].

RBF-based deformation is particularly attractive for adjoint-based optimization since it enables smooth, global geometry variations driven by a compact set of design variables and does not rely on mesh connectivity. This property makes the method robust for complex or unstructured meshes and well suited for repeated geometry updates in gradient-based aerodynamic optimization frameworks [23, 5].

4. Optimization Procedure

The aerodynamic shape optimization follows a standard gradient-based adjoint workflow, as commonly adopted in aerodynamic design optimization studies [29, 30, 5]. The process couples unsteady flow simulations, continuous adjoint sensitivity analysis, and geometry updates within an iterative loop reported in figure 1.

The airfoil geometry is parameterized by a finite set of control points distributed along the surface. The design variables β correspond to the vertical displacements of these points, allowing localized shape modifications on both suction and pressure sides while maintaining sufficient geometric flexibility. Similar surface-based parameterizations have been widely used in adjoint airfoil optimization to balance design freedom and numerical robustness [5]. Bound constraints are imposed to guarantee physically meaningful deformations, and a penalty formulation is employed to prevent surface overlap and enforce a minimum airfoil thickness, following the strategy proposed in [5].

For a given geometry, the unsteady incompressible flow is first solved to obtain the primal solution. Based on this solution and the selected objective functional, appropriate adjoint boundary conditions are imposed on the airfoil surface, consistent with continuous adjoint theory for incompressible flows [29, 30]. The adjoint equations are then solved, either in fully unsteady form or using a time-averaged primal state to reduce computational cost, an approach previously shown to be effective for aerodynamic objectives dominated by mean forces [5].

Once the adjoint solution is available, shape sensitivities with respect to the design parameters are evaluated using the continuous adjoint formulation. These gradients are subsequently used to update the design variables through a gradient-based optimization algorithm. In this work, a limited-memory BFGS (L-BFGS) method with bound constraints is employed, following established practice in large-scale aerodynamic optimization problems [31].

After each optimization step, the airfoil geometry is updated using the RBF deformation method, and the procedure is repeated until convergence criteria based on the objective value and/or gradient norm are satisfied. This adjoint-driven workflow enables efficient and robust aerodynamic optimization for unsteady turbulent flows, as discussed in [5, 24].

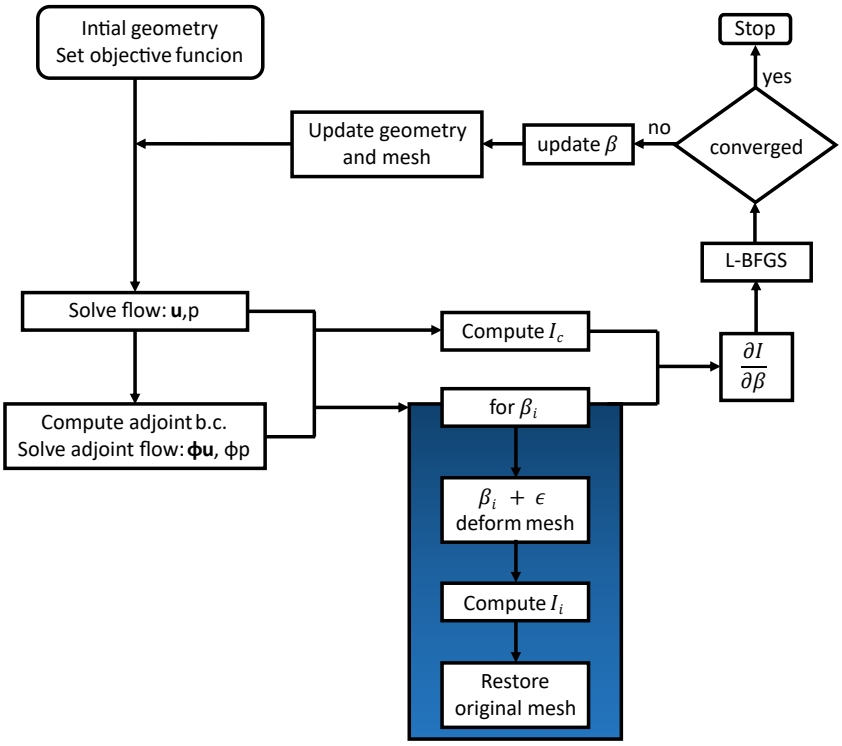


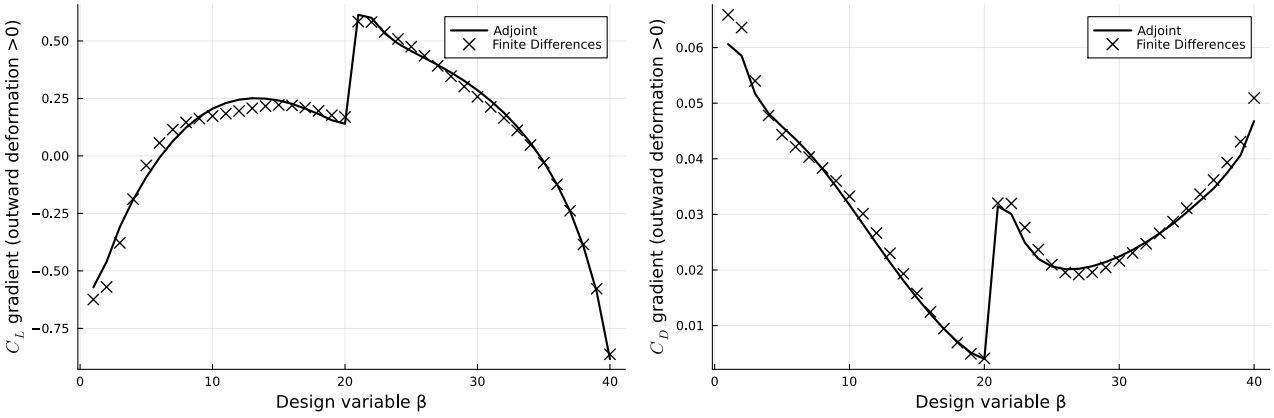
Figure 1 – Flow-chart detailing the steps of the adjoint-based airfoil surface aerodynamic optimization process

5. Finite Difference Validation

In this section the validation with FD results is presented. The case study is the flow over a NACA 0012 airfoil, with an Angle of Attack (AoA) of 2.5° at Reynolds number 1000. Quadrangular second order elements ($Q2$) are employed to ensure low numerical dissipation [32]. The *flow through time* (ftt) is introduced for normalization: $ftt = c/|\mathbf{u_0}|$, where c is the airfoil chord, and $\mathbf{u_0}$ is the inlet velocity. The time-step used is 0.05 and flow-field solutions are time-averaged for $5.0\,\mathrm{ftt}$, after discarding the first $4.0\,\mathrm{ftt}$ of simulations.

5.1 Validation of geometrical sensitivity computed using the adjoint solution

For the FD validation, each control point has been displaced vertically in the outward direction by 1×10^{-4}. In figure 2 the comparison of the gradient computed with the adjoint and FD for 2 different objective functions (C_L, C_D) are displayed, showing an excellent agreement between the two methods and therefore validating the implementation. The differences are negligible, the trend is correctly captured, and most importantly, the sign of the gradient is correct, meaning the optimizer correctly determines whether it needs to move the geometry inward or outward.



(a) C_L gradient

(b) C_D gradient

Figure 2 – Finite Difference validation, NACA 0012 airfoil, Reynolds number 1000, AoA 2.5°

Furthermore, the curves are very similar to the one obtained in [7, 33] for the same test case, but using different adjoint and deformation methods.

6. Inverse Geometrical Optimization

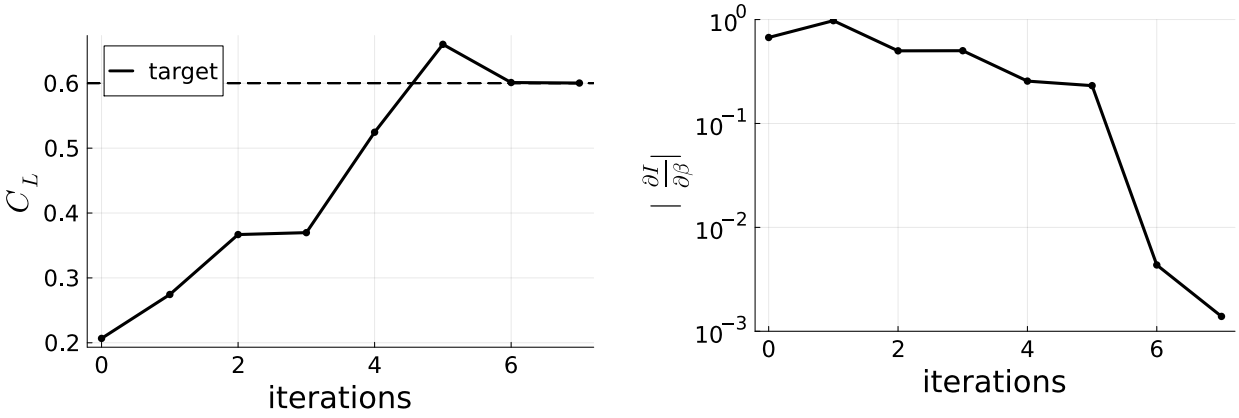
The first case to validate the adjoint implementation and workflow consists of finding an airfoil with: $\overline{C_L^*} = 0.6$ at $AoA = 4.0^\circ$ and $Re = 1000$. The NACA 0012 airfoil is used as the initial geometry, similarly as in [5]. The unsteadiness in the present test case arises from laminar vortex shedding at a low Reynolds number, well below the transition threshold, so the entire flowfield remains fully laminar throughout the shedding cycle. Consequently, the two dimensional Navier Stokes equations are sufficient to capture the relevant flow physics, and no turbulence closure is required. This distinguishes the present case from configurations involving a laminar separation bubble followed by turbulent reattachment, where the downstream flow would be inherently three dimensional and chaotic, and where the accuracy of the resulting gradients would be far more difficult to justify.

The primal and adjoint simulations are run for $10.0\,\mathrm{fft}$ and only the last $2.0\,\mathrm{fft}$ are kept to avoid the initial transient. The objective function is given $I_C = \frac{1}{2}(\overline{C_L^*} - \overline{C_L})^2$.

The optimization cycle is reported in figure 3 and the magnitude of the gradient of the objective function with respect to the shape parameters, $\left|\frac{\partial I}{\partial \beta}\right|$, drops by 3 orders of magnitude. The initial configuration exhibits a steady solution, whereas the optimized airfoil produces an unsteady flow.

7. Computational Cost and Comparison with RANS-Based Approaches

The primary computational cost of the proposed VMS-consistent unsteady adjoint framework arises from two sources: the need to resolve the full unsteady flow physics at a resolution sufficient for LES,



(a) C_L evolution with iterations

(b) Magnitude of the gradient evolution with iterations

Figure 3 – Optimization with target $\overline{C_L^*} = 0.6$, starting from NACA0012

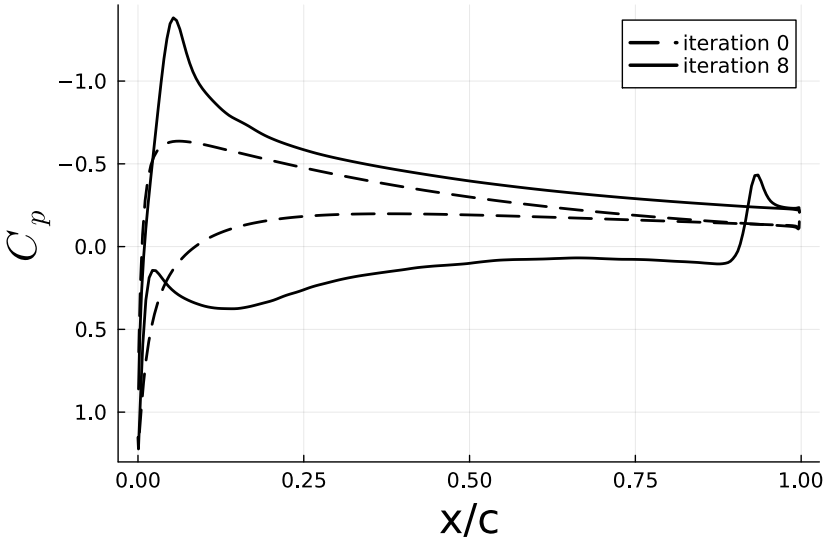


Figure 4 – Initial and final time-averaged C_p distribution, for a target $\overline{C_L^*} = 0.6$

and the requirement to store or recompute the primal solution over the entire optimization time horizon for backward-in-time adjoint integration. Compared to steady-state RANS-based adjoint methods, this results in a substantially higher computational and memory footprint per optimization iteration. The memory burden can be partially alleviated through checkpointing strategies such as the Revolve algorithm of Griewank and Walther [34]. RANS-based adjoint approaches benefit from low computational cost and mature implementations, but their applicability is limited in flow regimes where empirical turbulence closures introduce inaccurate predictions or non-differentiable behavior. A common practical issue is the frozen turbulence approximation, in which the turbulent eddy viscosity is held constant during adjoint assembly, leading to gradient errors that can compromise convergence or even the sign of the computed sensitivities [35, 14]. Even when the turbulence model is fully differentiated, RANS closures are known to be unreliable for flows involving LSBs, transitional boundary layers, or massively separated wakes, precisely the configurations where unsteady effects dominate the aerodynamic response. This limitation is most pronounced at low Reynolds numbers, where transition and separation phenomena are strongly sensitive to the modeling assumptions. At higher Reynolds numbers, where transition occurs earlier and boundary layers are more resistant to separation, RANS closures tend to recover a level of reliability closer to that observed for fully turbulent, attached flows.

LES of airfoil flows near stall, such as those reported by Frère et al. [36], reveal complex unsteady dynamics, including intermittent separation and vortex shedding, that RANS models fundamentally

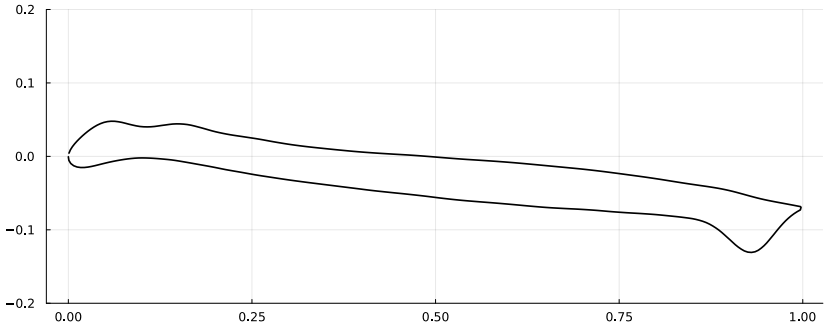


Figure 5 – Final shape resulting from adjoint optimization, for a target $\overline{C_L^*} = 0.6$

are not able to capture. In such regimes, RANS adjoint sensitivities may be unreliable or altogether unavailable, requiring ad hoc regularization strategies that compromise physical consistency.

The present VMS-LES adjoint formulation avoids these issues by deriving the adjoint equations directly from the same variational framework employed in the primal solver. This ensures that sensitivities accurately reflect the resolved-scale physics without reliance on empirical closures or differentiation approximations. For problems governed by unsteady phenomena such as vortex shedding, laminar-to-turbulent transition, or dynamic stall, the unsteady adjoint provides gradient information that is physically consistent and unobtainable from any steady-state analysis. The extension to unsteady primal flows introduces challenges beyond the memory cost, including the time-accurate backward integration of the adjoint equations and the correct treatment of time-periodic or statistically stationary objective functions.

Although computationally more demanding, the proposed methodology provides a favorable accuracy-to-cost ratio for problems where unsteady effects dominate the aerodynamic response and where RANS-based methods fail to deliver dependable gradients. For the two-dimensional, moderate-Reynolds-number configurations considered in this study, the overall computational cost remains tractable, demonstrating the feasibility of high-fidelity unsteady adjoint-based optimization.

Remarks on Computational Cost

The numerical experiments presented in this work are carried out on the NACA 0012 airfoil at $\alpha = 4^\circ$ and $Re = 1000$ in two dimensions, using a structured mesh of first-order quadrangular elements with $44\,287$ nodes, corresponding to $88\,574$ velocity degrees of freedom and $132\,861$ total degrees of freedom once pressure is included. The time step is fixed at $\Delta t = 0.05\,\mathrm{ftt}$, and the first $10.0\,\mathrm{ftt}$ of each simulation are discarded to allow the flow to develop, leaving 1000 time steps over which the adjoint is integrated backward in time.

The dominant cost driver of the present framework is the backward-in-time adjoint integration, which requires access to the full primal trajectory. Without compression, storing all 1000 velocity snapshots at $88\,574$ components each amounts to $1000 \times 88\,574 \times 8\,\mathrm{bytes} \approx 700\,\mathrm{MB}$ for double-precision velocity alone, a value that grows rapidly with the number of time steps, spatial resolution, dimensionality, and element order. In contrast, a steady RANS adjoint requires storing only a single primal flow field and its associated adjoint.

It should be noted that at $Re = 1000$, the flow over the NACA 0012 is predominantly laminar, and a steady laminar or low-Reynolds RANS computation can reproduce the mean flow with reasonable accuracy at a fraction of the cost. In this particular regime, the advantage of the unsteady adjoint lies not in correcting for turbulence modeling errors, but in capturing the sensitivity of time-averaged objectives to the unsteady laminar dynamics, including the periodic shedding and the LSB bursting that govern the aerodynamic response. The shortcomings of RANS-based adjoints become more severe as the Reynolds number increases toward fully turbulent or transitional regimes, where empirical closures and transition models introduce the inconsistencies discussed above [35, 14].

A further practical distinction concerns mesh resolution requirements. Because RANS models the entire turbulent energy cascade through empirical closure, the mesh needs only to resolve the mean flow and satisfy near-wall y^+ criteria, allowing relatively coarse grids. LES, by contrast, must re-

solve the energy-containing eddies directly, and the required mesh density grows substantially with Reynolds number. The total cost grows exponentially with the Reynolds number $\propto Re^{9/4}$ [37]. For the low-Reynolds-number two-dimensional case considered here, this additional resolution requirement remains manageable, but it represents a key cost factor as the framework is extended to higher Reynolds numbers or three-dimensional configurations.

8. Conclusions & Future Perspectives

A continuous unsteady adjoint method consistent with the Variational Multiscale formulation has been developed and applied to airfoil shape optimization in unsteady, low-Reynolds-number flows. The adjoint equations are derived directly from the stabilized VMS finite-element formulation used for the primal Large Eddy Simulation, guaranteeing strict numerical consistency between the primal and adjoint operators, including stabilization and subgrid-scale contributions.

The adjoint implementation has been rigorously validated through finite-difference sensitivity analysis for lift and drag on a NACA 0012 airfoil, showing excellent agreement in both magnitude and sign of the gradients. This validation confirms the correctness of the continuous VMS-adjoint formulation and its suitability for aerodynamic sensitivity analysis. The full adjoint-driven optimization framework has been successfully demonstrated through an inverse design problem targeting a prescribed time-averaged lift coefficient, achieving convergence and producing a physically consistent optimized geometry. The authors would like to emphasize that this work is still in progress, and the current findings should be viewed as preliminary results within a broader, ongoing research effort.

The proposed approach offers significant advantages over traditional RANS-based adjoint methods, particularly in flow regimes involving LSBs, transition, and unsteadiness, where empirical turbulence models and their adjoints often lack robustness or consistency. By embedding the adjoint directly within the VMS-LES framework, this methodology enables high-fidelity, physically accurate shape optimization for problems where steady or Reynolds-averaged assumptions are invalid. Although the unsteady adjoint entails a substantially higher computational and memory cost than its RANS counterpart, owing to the backward-in-time integration over the full primal trajectory and the finer mesh resolution required by scale-resolving simulations, this overhead is an intrinsic consequence of the higher-fidelity physical modeling and remains tractable for the two-dimensional configurations considered here. The present work therefore represents an important step toward robust, high-fidelity unsteady adjoint optimization for aerospace applications at practical Reynolds numbers, and motivates future extensions to three-dimensional, fully turbulent configurations where the limitations of RANS-based sensitivity methods are most acute.

Several directions for future work emerge naturally from the present study. The most immediate extension concerns the generalization of the framework to three-dimensional configurations, where the full complexity of turbulent structures and spanwise flow features can not be predict by two-dimensional simulations, and where the cost of the backward-in-time adjoint integration makes efficient checkpointing strategies essential. A central open question is the behavior of the adjoint in fully turbulent and transitional regimes at higher Reynolds numbers, where LSBs develop. In chaotic or fully turbulent regimes, nearby flow trajectories diverge exponentially, and the adjoint linearization inherits this instability, growing unboundedly as it is integrated backward over long time horizons [38]. The resulting gradient is then dominated by numerical noise rather than physical sensitivity information, rendering the standard adjoint unreliable as a descent direction for optimization. Exploring the stability limits of the continuous VMS adjoint in such regimes, and investigating dedicated sensitivity methods such as least-squares shadowing [38], represents a central open challenge that will guide the future development of the present framework.

9. Contact Author Email Address

Corresponding author: carlo.brunelli@mil.be
Authors: matija.avirovic@mil.be, bart.janssens@mil.be, benoit.marinus@mil.be

10. Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder

of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.

References

[1] F. Angrand. Optimum design for potential flows. *International Journal for Numerical Methods in Fluids*, 3:265–282, 5 1983.

[2] François Beux and Alain Dervieux. Exact-gradient shape optimization of a 2-d euler flow. *Finite Elements in Analysis and Design*, 12:281–302, 12 1992.

[3] W. Kyle Anderson and Daryl L. Bonhaus. Airfoil design on unstructured grids for turbulent flows. *AIAA Journal*, 37:185–191, 1999.

[4] W.Kyle Anderson and V. Venkatakrishnan. Aerodynamic design optimization on unstructured grids with a continuous adjoint formulation. *Computers & Fluids*, 28:443–480, 5 1999.

[5] D. N. Srinath and S. Mittal. A stabilized finite element method for shape optimization in low reynolds number flows. *International Journal for Numerical Methods in Fluids*, 54:1451–1471, 8 2007.

[6] M Schramm, B Stoevesandt, and J Peinke. Adjoint Optimization of 2D-Airfoils in Incompressible Flows. In *11th World Congress on Computational Mechanics*, 2014.

[7] Carlos Castro, Carlos Lozano, Francisco Palacios, and Enrique Zuazua. Systematic continuous adjoint approach to viscous aerodynamic design on unstructured grids. *AIAA Journal*, 45:2125–2139, 9 2007.

[8] Karthik Mani and Dimitri Mavriplis. Geometry Optimization in Three-Dimensional Unsteady Flow Problems using the Discrete Adjoint. In *51st AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, Grapevine (Dallas/Ft. Worth Region), Texas, 2013. American Institute of Aeronautics and Astronautics.

[9] Siva K. Nadarajah and Antony Jameson. Optimum Shape Design for Unsteady Flows with Time-Accurate Continuous and Discrete Adjoint Method. *AIAA Journal*, 45(7):1478–1491, 2007.

[10] Markus P. Rumpfkeil and David W. Zingg. The optimal control of unsteady flows with a discrete adjoint method. *Optimization and Engineering*, 11(1):5–22, 2010.

[11] Kyriakos C. Giannakoglou and Dimitrios I. Papadimitriou. Adjoint Methods for Shape Optimization. In Dominique Thévenin and Gábor Janiga, editors, *Optimization and Computational Fluid Dynamics*, pages 79–108. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[12] D. N. Srinath and Sanjay Mittal. An adjoint method for shape optimization in unsteady viscous flows. *Journal of Computational Physics*, 229:1994–2008, 3 2010.

[13] Qiqi Wang and Jun Hui Gao. The drag-adjoint field of a circular cylinder wake at reynolds numbers 20, 100 and 500. *Journal of Fluid Mechanics*, 730:145–161, 2013.

[14] A.S. Zymaris, D.I. Papadimitriou, K.C. Giannakoglou, and C. Othmer. Continuous adjoint approach to the spalart–allmaras turbulence model for incompressible flows. *Computers & Fluids*, 38:1528–1538, 9 2009.

[15] Michael G. H. Piotrowski and David W. Zingg. Smooth Local Correlation-Based Transition Model for the Spalart–Allmaras Turbulence Model. *AIAA Journal*, 59(2):474–492, 2021.

[16] Carlo Brunelli, Matija Avirovic, Bart Janssens, Benoît G. Marinus, Georg May, and Mark Runacres. Prediction of a Laminar Separation Bubble on Low-Reynolds Airfoils Using Variational Multiscale Method. In *AIAA AVIATION FORUM AND ASCEND 2024*, Las Vegas, Nevada, July 2024. American Institute of Aeronautics and Astronautics.

[17] Carlo Brunelli, Matija Avirovic, Bart Janssens, Benoît G Marinus, Mark Runacres, Georg May, Antoine Ecorce, and Antonin Dalbera. Passive flow control devices on low-Reynolds DU89-134 airfoil. In *34th Congress of the International Council of the Aeronautical Sciences*, 2024.

[18] Thomas J.R. Hughes, Gonzalo R. Feijóo, Luca Mazzei, and Jean-Baptiste Quincy. The variational multiscale method—a paradigm for computational mechanics. *Computer Methods in Applied Mechanics and Engineering*, 166:3–24, 11 1998.

[19] Thomas J R Hughes. Green's functions, the dirichlet-to-neumann formulation, subgrid scale models, bubbles and the origins of stabilized methods. *Computer Methods in Applied Mechanics and Engineering*, 127:387–401, 1995.

[20] Thomas J.R. Hughes, Luca Mazzei, and Kenneth E. Jansen. Large eddy simulation and the variational multiscale method. *Computing and Visualization in Science*, 3:47–59, 5 2000.

[21] Y. Bazilevs, V. M. Calo, J. A. Cottrell, T. J.R. Hughes, A. Reali, and G. Scovazzi. Variational multiscale residual-based turbulence modeling for large eddy simulation of incompressible flows. *Computer Methods in Applied Mechanics and Engineering*, 197:173–201, 12 2007.

[22] Bart Janssens, Pierre Vandenschrick, Kasper Stevens, and Giacomo Alessi. The continuous adjoint approach applied to the stabilized finite-element formulation of the incompressible Navier-Stokes equations. In *13th European Conference on Turbomachinery Fluid Dynamics & Thermodynamics*, 2019.

[23] A. de Boer, M. S. van der Schoot, and H. Bijl. Mesh deformation based on radial basis function interpolation. *Computers and Structures*, 85(11-14):784–795, June 2007.

[24] Jianping Niu, Juanmian Lei, and Jiandong He. Radial basis function mesh deformation based on dynamic control points. *Aerospace Science and Technology*, 64:122–132, May 2017.

[25] Carlo Brunelli. SegregatedVMSSolver.jl: Linearized and Segregated Stabilized Solver for Large Eddy Simulation in Julia. *Journal of Open Source Software*, 10(110):7564, June 2025.

[26] Carlo Brunelli, Matija Avirovic, Bart Janssens, Koen Hillewaert, Benoît G. Marinus, and Mark Runacres. A Linearized and Segregated Variational Multiscale Method for Incompressible Flows. *International Journal for Numerical Methods in Fluids*, page fld.70062, 2026.

[27] Carlo Brunelli, Matija Avirovic, Bart Janssens, Benoît G. Marinus, Koen Hillewaert, and Mark Runacres. Prediction of Laminar Separation Bubble on Airfoils at Low Reynolds Number. *Flow, Turbulence and Combustion*, 116(2):27, 2026.

[28] Matija Avirović, Carlo Brunelli, Benoît G. Marinus, Joris Degroote, and Jeroen Van Beeck. Laminar–Turbulent Transition on a DU89-134/14 Airfoil at Low Reynolds Number. *AIAA Journal*, pages 1–20, April 2026.

[29] O. Soto and R. Löhner. Cfd shape optimization using an incomplete-gradient adjoint formulation. *International Journal for Numerical Methods in Engineering*, 51:735–753, 6 2001.

[30] Orlando Soto and Rainald Lohner. On the Boundary Computation of Flow Sensitivities. In AIAA Aerospace Sciences Meeting and Exhibit, editor, *42nd AIAA Aerospace Sciences Meeting and Exhibit*, Reston, Virigina, January 2004. American Institute of Aeronautics and Astronautics.

[31] Richard H. Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16:1190–1208, 9 1995.

[32] Carlo Brunelli. SegregatedVMSSolver.jl: Linearized and Segregated Stabilized Solver for Large Eddy Simulation in Julia. *Journal of Open Source Software*, 10(110):7564, June 2025.

[33] Giuseppe Sorgiovanni, Maurizio Quadrio, and Raffaele Ponzini. A robust open-source adjoint optimization method for external aerodynamics. 2016.

[34] Andreas Griewank and Andrea Walther. Algorithm 799: Revolve: An Implementation of Checkpointing for the Reverse or Adjoint Mode of Computational Differentiation. *ACM Transactions on Mathematical Software*, 26:19–45, 3 2000.

[35] Richard P. Dwight and Joel Brezillon. Effect of Approximations of the Discrete Adjoint on Gradient-Based Optimization. *AIAA Journal*, 44(12):3022–3031, 2006.

[36] A. Frère, K. Hillewaert, P. Chatelain, and G. Winckelmans. High Reynolds Number Airfoil: From Wall-Resolved to Wall-Modeled LES. *Flow, Turbulence and Combustion*, 101(2):457–476, 2018.

[37] Ugo Piomelli. Wall-layer models for large-eddy simulations. *Progress in Aerospace Sciences*, 44(6):437–446, August 2008.

[38] Qiqi Wang, Rui Hu, and Patrick Blonigan. Least Squares Shadowing sensitivity analysis of chaotic limit cycle oscillations. *Journal of Computational Physics*, 267:210–224, 2014.