

Adjoint Formulation Investigations of Benchmark Aerodynamic Design Cases in SU2

Thomas D. Economon*, Juan J. Alonso[†]
Stanford University, Stanford, CA, 94305, U.S.A.

Timothy Albring[‡] and Nicolas R. Gauger[§]
Technical University of Kaiserslautern, Kaiserslautern, 67663, Germany

We investigate multiple adjoint formulations within the SU2 suite: a continuous adjoint in surface integral form, a continuous adjoint in field integral form, and a discrete adjoint leveraging algorithmic differentiation. The formulations are applied to the computation of drag sensitivities with respect to shape changes for four different external flow test cases, including the NACA 0012, RAE 2822, NACA 4412, and the Ahmed car body configuration. The main objective of this article is to draw meaningful conclusions about the performance of the available adjoint technologies for computing shape design gradients with a particular emphasis on accuracy. While all formulations provide gradients suitable for optimization purposes, we find that the inclusion of mesh effects in the field integral formulation of the continuous adjoint can provide improvements in sensitivity accuracy and robustness over the surface integral formulation.

I. Introduction

Adjoint methods for use in gradient-based optimal shape design have come a long way since early pioneering efforts in incompressible flows and airfoil design.^{1,2} Today, these methods are applied to large-scale problems of interest in the aircraft, automotive, and other engineering industries, and their effectiveness has been well documented in a rich body of literature over the past several decades for aerodynamic design in smooth, attached flows.

Due to the currently prohibitive cost of including scale-resolving simulations (Detached Eddy Simulation, Large Eddy Simulation, Direct Numerical Simulation) within the industrial design cycle in an automatic way, it is expected that Reynolds-averaged Navier-Stokes (RANS) will be the workhorse equation system for engineering design for years still to come. Given this situation, steady RANS is being pushed to its limits when embedded in shape design processes by calculating complex flows on highly-stretched, unstructured meshes composed of hundreds of millions of grid cells that may feature flow-induced or geometrically-induced separation. In light of this, it is imperative to understand the limitations of, and any opportunities for, the available adjoint-based techniques when computing aerodynamic sensitivities in these complicated situations with RANS.

Adjoint formulations are typically classified by their derivation as either continuous (the governing equations are first linearized then the result is discretized) or discrete (the governing equations are first discretized and the result is linearized). The merits for each approach have been discussed at length in literature, and they have been compared directly by several authors.³⁻⁵ While it is generally agreed that both approaches are equally useful sensitivity analysis techniques for aerodynamic shape design, the details of their implementation result in each having advantages and disadvantages under different circumstances. Even within the major categories of continuous and discrete, there are many different approaches to formulating the problem.

*Postdoctoral Scholar, Department of Aeronautics & Astronautics, AIAA Senior Member.

[†]Professor, Department of Aeronautics & Astronautics, AIAA Associate Fellow.

[‡]Ph.D. Candidate, Department of Mathematics, AIAA Student Member.

[§]Professor, Department of Mathematics, AIAA Senior Member.

For example, in recent years, practitioners of the continuous adjoint have been reconsidering the effects of the computational grid on the accuracy of shape sensitivities. From a mathematical point of view, the continuous adjoint equations and sensitivity formulas can be derived without any consideration given to discretization choices or even the existence of a grid on which to solve the equations. Some early derivations of the continuous adjoint did maintain terms related to the impact of grid displacements on the sensitivities in the form of field integrals.⁶ Later derivations removed the field integral terms entirely, resulting in sensitivity formulas that are surface integrals alone.^{7–9} The surface formulation is elegant, exact, and efficient: there is no need to perturb the grid, compute field integrals, or consider grid sensitivities.

It is important to note that these two continuous formulations based on surface and field integrals are mathematically equivalent, as demonstrated by Lozano et al.¹⁰ But, in practice, they can provide very different results. This was first noted by Anderson and Venkatakrishnan¹¹ as inaccuracies arising from higher-order derivative terms appearing in the surface formulation for unstructured grids. Castro et al.⁸ offered a new surface formulation that alleviated these issues and improved sensitivity accuracy by focusing on simplifications and introducing shape calculus. Lozano¹² presented clear discrepancies when comparing a surface formulation to finite differences for inviscid problems, and it was attributed to the fact that some analytic manipulations in the derivation do not hold at the discrete level, especially near geometric singularities (e.g., trailing edges). For turbulent flows, Lozano et al.¹⁰ concluded that the field integral formulation provides improvements in accuracy over the surface integral approach, but that the cost can be prohibitive unless the field integral terms are treated in a way that removes dependence on the number of design variables. Recently, Kavvadias et al.¹³ have presented a detailed discussion of the differences between the surface and field formulations and have also proposed a purely continuous method for including grid sensitivity terms that does not scale linearly with the number of design variables.

Clearly, there are many intricacies in the development of the various adjoint techniques, and it is often difficult to perform direct comparisons of the methods and to fairly assess their performance. However, thanks to collaboration across the open-source community, the SU2 software suite^{14–16} contains three distinctly different implementations of the adjoint methodology all within the same codebase: a continuous adjoint in a surface integral (SI) formulation, a continuous adjoint in a field integral formulation (FI), and a discrete adjoint constructed via algorithmic differentiation (AD) of the code. When combined together, these complementary sets of adjoint formulations offer an unprecedented opportunity to study a breadth of methods within a single investigation: continuous and discrete, different treatments of adjoint turbulence, and field integral and surface integral formulations.

With that in mind, the main objective of this paper is to draw meaningful conclusions about the performance of the available adjoint technologies for computing shape design gradients with a particular emphasis on accuracy. To accomplish this, we present the systematic application of a variety of adjoint technologies for computing aerodynamic sensitivities for a suite of test cases. For all investigations, we will look at sensitivities of drag with respect to shape changes of the specific geometry. In particular, we will focus on the differences between SI and FI for continuous formulations and also on the differences between frozen and differentiated turbulence models with the discrete adjoint.

This paper is organized as follows. Section II details the different adjoint formulations that are available in SU2 at a high-level. Then, in Section III, the series of test cases are studied one-by-one. A simple inviscid NACA 0012 is the first case. Second, we study the RAE 2822 airfoil in turbulent flow. The third case is the NACA 4412 airfoil at an angle of attack sufficient to cause mild flow-induced separation near the trailing edge of the upper surface. Finally, we explore the 3D Ahmed body geometry, which acts as a representative configuration for automotive applications. The Ahmed case features a large separation region in its wake induced by the sharp corners at the rear of the body, and it represents a more challenging case. Finally, in Section IV, we will summarize the major conclusions of the investigation.

II. Adjoint Formulations within SU2

In this section, we present a high-level description of the three different adjoint formulations that we will investigate. Here, we will focus on the final sensitivity expressions for the different formulations in order to highlight the key differences, and we refer the reader to the cited literature for full details. In addition to these adjoint implementations, SU2 contains all of the tools necessary for aerodynamic shape design, including the modules for solving primal and adjoint problems, acquiring gradient information by projecting sensitivities into the design space, and mesh deformation techniques. Typically, geometry parameterization

is accomplished with either Hicks-Henne bumps in 2D¹⁷ or a free-form deformation (FFD) approach in 2D or 3D,¹⁸ mesh deformation with an elasticity-based approach,¹⁹ and the shape design loop is driven by the SciPy SLSQP optimizer.²⁰

A. Continuous Adjoint in a Surface Integral Formulation (SI)

SU2 features a native continuous adjoint solver in a surface integral formulation based on the frozen viscosity, or frozen turbulence (FT), assumption.⁸ A continuous adjoint that includes linearization of the Spalart-Allmaras (S-A) turbulence model has also been developed.²¹ In the surface formulation, the final expression for the variation of a typical force-based objective, e.g., drag or lift, with respect to infinitesimal surface perturbations δS can be written as

$$\delta \mathcal{J} = \int_S \left(\vec{n} \cdot \bar{\bar{\Sigma}}^\varphi \cdot \partial_n \vec{v} - \mu_{tot}^* c_p \nabla_S \psi_{\rho E} \cdot \nabla_S T \right) \delta S ds, \quad (1)$$

where \mathcal{J} is the chosen objective, \vec{n} is the local surface normal, $\bar{\bar{\Sigma}}^\varphi$ is the adjoint stress tensor, v is the primal velocity in a Cartesian reference frame, $\mu_{tot}^* = \frac{\mu_{dyn}}{Pr_d} + \frac{\mu_{tur}}{Pr_t}$ is the effective thermal conductivity, c_p is the ratio of specific heats at constant pressure, $\psi_{\rho E}$ is the adjoint energy, T is the temperature, $\partial_n(\cdot) = \vec{n} \cdot \nabla(\cdot)$ is the normal gradient operator at a surface point, and $\nabla_S(\cdot) = \nabla(\cdot) - \partial_n(\cdot)\vec{n}$ is the tangential gradient operator at a surface point. Note that Eqn. 1 is a *surface* integral and that there are no field integral contributions to the sensitivity or any terms related to the volume grid. This formulation allows for the visualization of surface sensitivity maps by evaluating the expression in parentheses inside the integral in Eqn. 1 at each grid point on the design surface.

Since debuting in the initial public release of SU2, the continuous adjoint solver has been extensively used and rigorously tested^{22,23} for both inviscid and viscous problems across many flow regimes. Although it does not provide numerically exact gradient information, it has been demonstrated as a good choice for aerodynamic shape design due to sufficient accuracy, its efficiency, and the ability to customize numerical methods for solving the adjoint equations. This method requires a single primal PDE solve and a single adjoint PDE solve to obtain the sensitivity with respect to an arbitrary number N of design variables.

B. Continuous Adjoint in a Field Integral Formulation (FI)

A continuous adjoint in a field integral formulation has been recently implemented within SU2, which offers another avenue for investigation. The formulation is based on the original adjoint work of Jameson as described in Leoviriyakit et al.,⁶ and the final sensitivity formula can be expressed as

$$\delta \mathcal{J} = \left(\frac{\partial \mathcal{J}}{\partial \alpha} - \int_\Omega \Psi^\top \frac{\partial \mathcal{R}}{\partial \alpha} d\Omega \right) \delta \alpha, \quad (2)$$

where \mathcal{J} is the chosen objective, α is the design variable, Ψ is the adjoint state vector, and \mathcal{R} is the system of governing equations written in residual form. The first term on the right-hand side is a partial derivative of the objective with respect to the design variable(s). The second term on the right-hand side is a field integral that requires the adjoint state at each node in the volume grid along with a partial derivative of the residual with respect to the design variable(s). Note that the partial derivatives here imply a fixed flow solution, meaning that they represent the changes in the objective and residual due solely to geometry in the case of shape design.

In practice, the FI formulation can be easily implemented alongside an existing SI continuous adjoint formulation. For example, assume we are interested in a drag objective for flow around an airfoil with a bump function that modifies the shape of the airfoil. First, we compute the primal and adjoint states on the original mesh and store the value of drag. Then, we impose a small perturbation of the bump that deforms not only the airfoil shape, but propagates deformations into the volume grid, which is typically accomplished via a pseudo-structural mesh deformation algorithm. The two partial derivatives in Eqn. 2 can then be computed with finite differencing by evaluating the drag with the fixed flow solution on the perturbed airfoil surface and also evaluating the residual on the perturbed volume grid with the fixed flow solution.

While the operations to compute the final sensitivity are very cheap individually and only require a single flow and adjoint solve, the cost of this method scales linearly with the number of design variables N , as the surface and volume grid must be perturbed to evaluate the sensitivity of each design variable.

However, if one can apply an analytic or otherwise cheap mesh deformation algorithm, this linear increase in cost is manageable. Note that we solve identical primal and adjoint PDE systems for both the SI and FI formulations with the only difference being the evaluation of the sensitivities in Eqns. 1 and 2.

C. Discrete Adjoint based on Algorithmic Differentiation (AD)

A discrete adjoint solver in SU2 has been implemented by algorithmically differentiating the entire codebase in combination with all of the existing design parameterizations available in SU2. To be more precise, this includes differentiating not only the flux routines, but also the complete nonlinear iteration structure of the solver. By exploiting the fixed-point structure, one can construct a duality-preserving iteration for solving the adjoint system *without* the need to manually construct an exact Jacobian matrix. In this approach, fixed-point problems are solved for the flow \bar{U} and grid \bar{X} discrete adjoint states that leverage algorithmic differentiation to “tape” a representation of the Jacobian matrix that can be evaluated. The final sensitivity formula reduces to

$$\frac{d\mathcal{J}}{d\alpha} = \frac{d}{d\alpha} M^T(\alpha) \bar{X}, \quad (3)$$

where \mathcal{J} is the chosen objective, α is the design variable vector, X represents the grid, and M is the system of equations handling the deformation of the surface and volume grid, i.e., $M(\alpha) = X$. This method requires a single primal PDE solve, an adjoint solve for the flow, and an adjoint solve for the grid sensitivities in order to obtain the total sensitivity with respect to an arbitrary number N of design variables.²⁴

To improve efficiency in terms of compute and memory overhead, the group at the Technical University of Kaiserslautern has developed a custom set of AD tools specifically tuned for CFD applications as an open-source C++ library *Code-Differentiation Package* (CoDiPack, <https://github.com/SciCompKL/CoDiPack>). The discrete adjoint solver has been implemented in a general way within the SU2 suite, such that discrete adjoints for any evolutions and extensions of the primal analysis capabilities will be immediately available upon compilation without the need to manually differentiate any new components for inclusion in the discrete adjoint.²⁵

In addition to the three adjoint formulations mentioned above, we will compare all sensitivity results against a Finite Differences (FD) approach. In practice, this is accomplished by solving the primal equations in the baseline configuration and then once for a small perturbation of each design variable to get the difference in the objective \mathcal{J} before dividing by the perturbation size $\Delta\alpha$. This process requires a total of $N + 1$ primal calculations and N mesh perturbations. For all FD results in this paper, $\Delta\alpha = 1e-5$.

It is important to note that each of the above formulations is a valid approach for performing sensitivity analysis, and all have been used successfully for aerodynamic design. However, it is no surprise that each has its own advantages and disadvantages. For example, the continuous adjoint is typically more efficient in terms of memory and compute than the discrete adjoint, but the discrete adjoint offers numerically exact gradient information while the continuous adjoint does not. The focus of this paper is not on declaring a single best approach. Rather, the focus is on using all of the available methods in concert in order to draw new conclusions and to obtain a deeper understanding of how the formulations perform.

III. Benchmark Cases

In this section, we investigate the following series of test cases: the NACA 0012 airfoil in inviscid, transonic flow, the RAE 2822 airfoil in turbulent, transonic flow, the NACA 4412 airfoil in turbulent, subsonic flow, and the Ahmed body configuration in turbulent, subsonic flow. For all cases, we will focus on the sensitivities of drag with respect to shape changes of the specific geometry, which will be computed with the continuous adjoint in surface integral formulation (SI), continuous adjoint in field integral formulation (FI), discrete adjoint based on algorithmic differentiation (AD), and finite differences (FD). For viscous flows, the frozen viscosity assumption is always applied for the continuous adjoint methods (SI FT and FI FT) but optionally applied for the discrete adjoint (AD FT when applied). The zero sensitivity line is also plotted on the figures in the following sections to help identify any differences in signs.

A. NACA 0012


The first case is the NACA 0012 in transonic, inviscid flow. The Mach number is 0.85, and the airfoil is at a zero angle-of-attack. Due to the zero angle of attack, a symmetry boundary condition can be used along the upstream and downstream axes from the leading and trailing edges of the airfoil, respectively, which allows for reduced computational cost by simulating only the upper half of the domain. Consequently, drag is reported only for the upper surface of the airfoil, rather than doubling the value for the full airfoil drag.

A family of regular (curvilinear) O-grids was generated for performing a grid refinement study to assess the behavior of the various adjoint formulations. Information about the grid family and primal drag values in counts is reported in Tab. 1. The 2D grid sizes are given as $N_x \times N_y$, where N_x is the number of grid points along the upper airfoil surface and far-field boundary and N_y is the number of grid points along the upstream and downstream boundaries along the symmetry plane. N is the total number of grid points in each mesh, and h is the non-dimensional grid spacing. Starting with the finest mesh, we remove every other grid line with each successive coarsening, resulting in a doubling of the grid spacing on each level. The far-field boundary is located approximately 150 chord lengths from the airfoil. Characteristic-based far-field boundary conditions are enforced on the outer boundary, and a flow tangency boundary condition is enforced on the airfoil.

Grid	$N_x \times N_y$	N	h	C_d
1	0065 x 0129	8385	16	236.42
2	0129 x 0257	33153	8	235.54
3	0257 x 0513	131841	4	235.51
4	0513 x 1025	525825	2	235.54
5	1025 x 2049	2100225	1	235.55

Table 1. The family of grids used for the NACA 0012 investigation and their drag values in counts.

Fig. 1(a) contains a view of the grid near the airfoil surface. The primal flow is computed using the JST²⁶ convective scheme. A modified JST scheme featuring only a higher-order dissipation term is applied to the convective flux in the continuous adjoint problem. Both the primal and continuous adjoint solves take advantage of nonlinear multigrid. Implicit, local time-stepping is used to converge the problem to the steady-state solution, and the linear system is solved using GMRES.²⁷ Fig. 1(b) displays the pressure contours on the finest mesh, which show the strong shock present on the upper surface for these conditions.



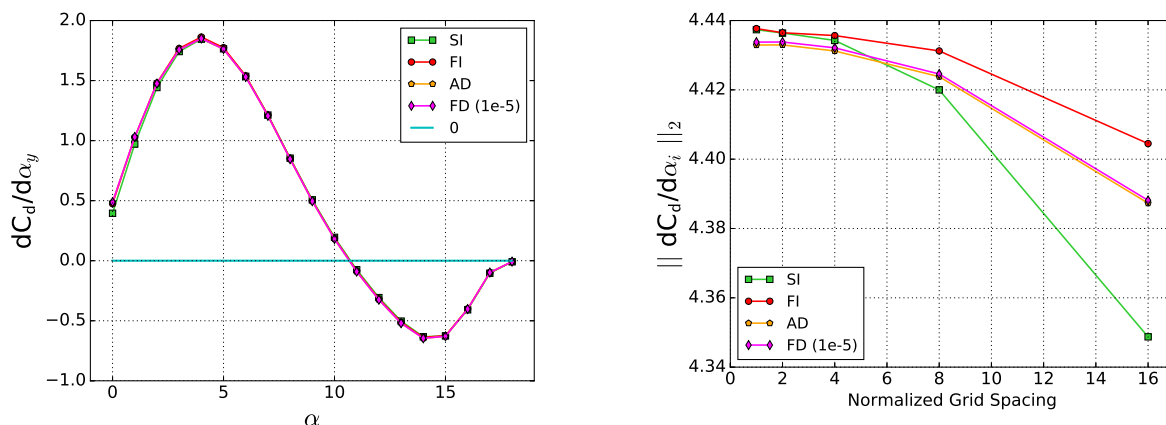
(a) View of the coarsest mesh near the airfoil surface.

(b) Coefficient of pressure contours near the airfoil surface on the finest mesh.

Figure 1. NACA 0012 pressure contours and a view of the mesh.

For this case, we will compare the drag sensitivity on the upper airfoil surface with respect to 19 Hicks-Henne bump variables that are equally spaced in 5% chord increments with the first located at 5% chord. Fig. 2(a) shows the drag gradient. The results from all formulations are in excellent agreement, even on the

coarsest mesh, with the only noticeable difference being a small discrepancy for the surface formulation near the leading edge.



(a) Drag sensitivities for 19 Hicks-Henne bumps on the coarsest mesh.

(b) 2-norm of the drag gradient vector as computed across the family of grids.

Figure 2. NACA 0012 drag sensitivities and behavior of the drag gradient during grid refinement.

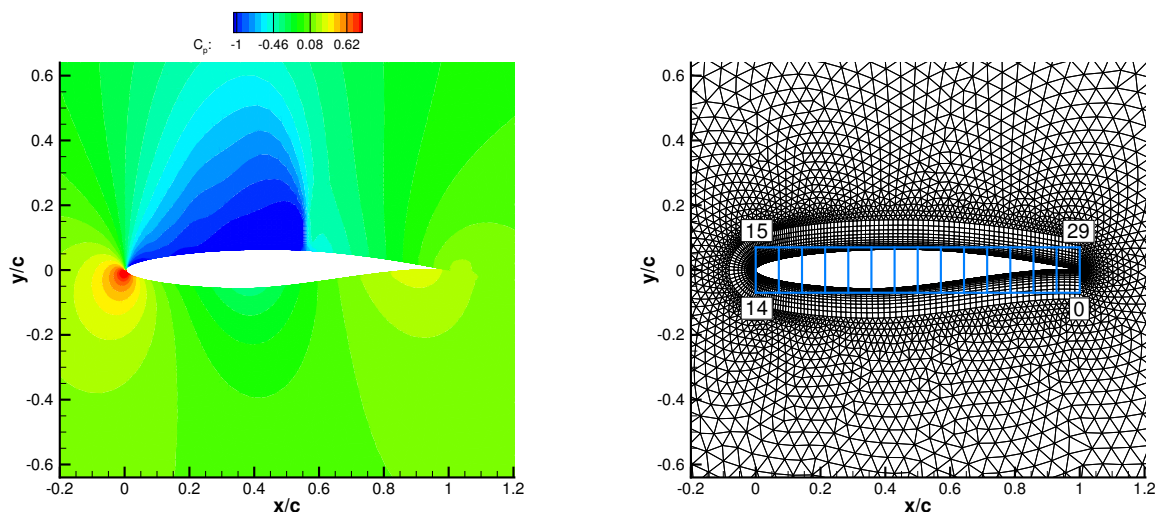
However, using our family of grids, we can investigate the behavior of the adjoint formulations in more detail. Fig. 2(b) presents the 2-norm of the drag gradient as we refine the grid while holding all of the numerical methods and design variable parameterization fixed. As expected, the various formulations all converge toward a zero grid-spacing value. The AD and FD results track each other very closely in this study and are nearly indistinguishable. However, the SI and FI formulations depart from the AD and FD results, as well as from each other. The continuous and discrete formulations converge to different values for the gradient magnitude as the grid is refined. This can be explained by the fact that they solve slightly different problems by construction in the present implementations when considering their linearization and discretization. We also see in Fig. 2(b) that, compared to each other, the FI gradient converges more quickly to the asymptotic value than the SI formulation, suggesting that including the field integral terms can reduce errors in the continuous formulation more quickly.

B. RAE 2822

The second case we will consider is the RAE 2822 airfoil in transonic, turbulent flow, which requires the solution of the RANS equations. The RAE 2822 airfoil is a supercritical airfoil commonly used for validation of turbulence models. Experimental case 6 from AGARD²⁸ is considered with flow conditions as described by NASA's NPARC Alliance Verification and Validation Archive.²⁹ The Mach number is 0.729, the angle of attack is 2.31° , the Reynolds number based on a chord of 1.0 is $6.5e6$, and the free-stream temperature is 288.15 K. For more SU2 validation results with the RAE 2822 case, the reader is referred to previous work.¹⁵ Fig. 3(a) presents pressure contours around the RAE 2822 for case 6.

We use a mixed-element grid, featuring a region of quadrilaterals wrapped around the near-body region of the RAE 2822 airfoil, similar to a C-grid topology for a structured mesh. It contains 13,937 grid points and 22,842 elements in total with 192 edges making up the airfoil boundary and 40 edges along the far-field boundary. The first grid point off the airfoil surface is at a distance to ensure $y^+ \approx 1$, and the far-field boundary is located approximately 100 chord lengths away from the airfoil. A characteristic-based far-field boundary condition is enforced on the far-field boundary, and a no-slip, adiabatic boundary condition is enforced on the airfoil. 30 FFD control point variables are defined above and below the airfoil. Both the mesh and FFD control points are shown Fig. 3(b).

For the primal problem, convective fluxes are computed with the Roe scheme³⁰ with a second-order reconstruction and the Venkatakrishnan limiter³¹ applied. Turbulent variables for the S-A model are convected with a first-order scalar upwind method, and viscous fluxes are calculated using the corrected average-gradient method. Implicit, local time-stepping is used to converge the problem to the steady-state solution, and the linear system is solved using GMRES. For the continuous adjoint problem, we use a second-order

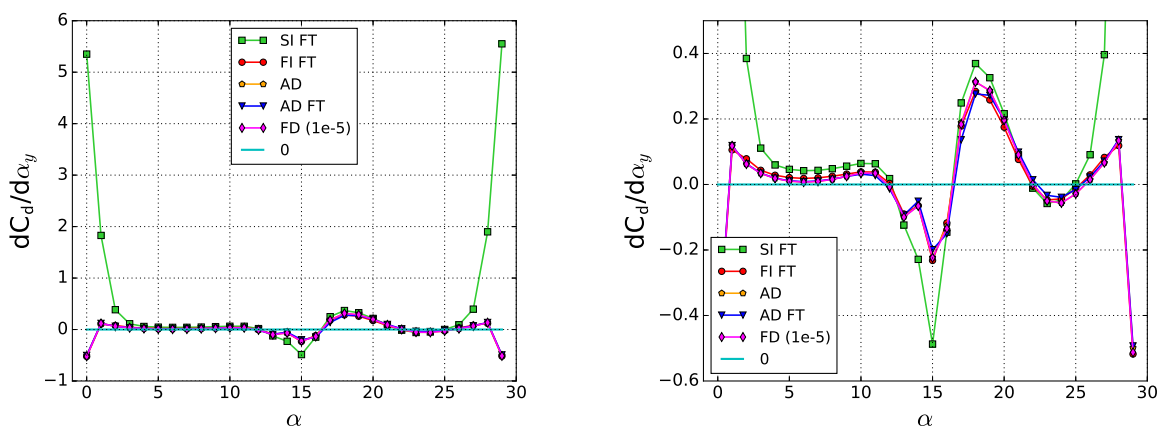


(a) Pressure contours around the RAE 2822 with the typical shock seen in Case 6. (b) View of the mesh near the airfoil surface with FFD control point numbering.

Figure 3. RAE 2822 pressure contours and a view of the grid near the airfoil.

upwind convective scheme and apply the frozen viscosity assumption, i.e., no adjoint turbulence model. We do not apply a slope limiter to the continuous adjoint solution.

Fig. 4(a) and 4(b) show full and zoom views, respectively, of the computed drag sensitivities with respect to vertical motion of the 30 FFD control variables. Overall, the gradients follow the same trends, and the signs of the sensitivity are mostly in agreement. However, there are several key differences to highlight. First, the SI FT sensitivities stray far from the others near the leading and trailing edges (sign and magnitude). This phenomenon was also noted by Lozano¹² and attributed to the presence of sharp edges, as the SI FT formulation assumes a smooth surface in its derivation. Alternatively, we see that the FI FT result closely matches the AD and FD results. The inclusion of the grid effects in the FI FT formulation not only corrects the issues at the sharp edges, but also moves the sensitivities closer to the AD and FD results along the smooth portions of the upper and lower surfaces, as seen in Fig. 4(b).



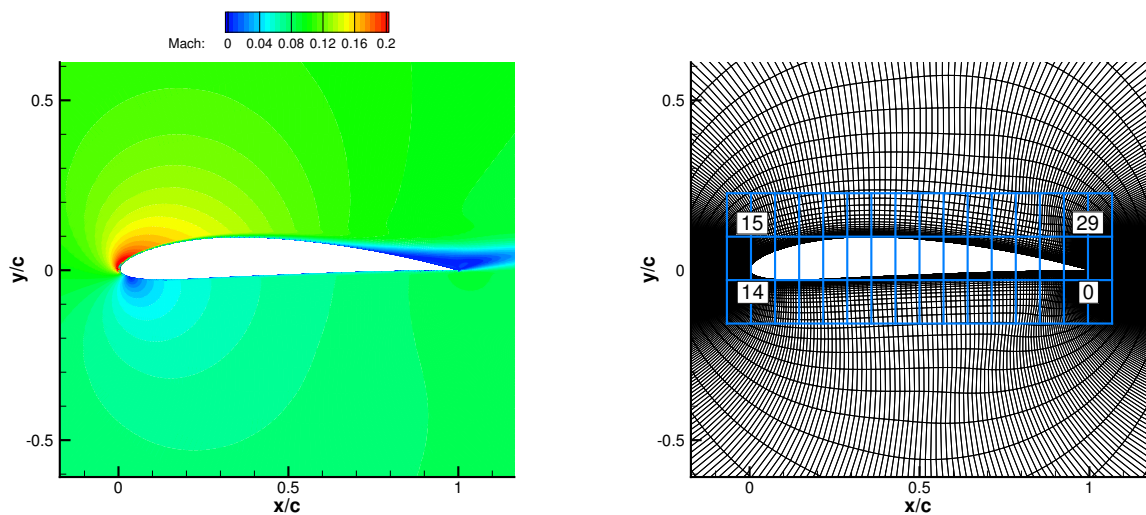
(a) Drag sensitivities with respect to vertical (y-axis) movement of 30 FFD control point variables.

(b) Zoom view of the drag sensitivities.

Figure 4. RAE 2822 drag sensitivity results.

The AD and FD results are nearly coincident, as expected. For this case, we also compute the discrete adjoint with frozen viscosity to find the relative importance of including an adjoint turbulence model for

this case. The AD FT result closely resembles the AD counterpart. Small discrepancies between AD and AD FT can be seen along the upper surface of the airfoil upstream and downstream of the shock, which is not surprising, as the RAE 2822 case 6 features a small shock-boundary layer interaction on the upper surface. However, the overall impact of freezing the turbulence in this case is relatively small, especially when compared to the impact of including grid effects with the continuous adjoint.



(a) Mach number contours. A small separation region is visible near the trailing edge on the upper surface. (b) View of the mesh near the airfoil surface with FFD control point numbering.

Figure 5. NACA 4412 Mach number contours and view of the grid.

C. NACA 4412

This case simulates turbulent flow over a NACA 4412 airfoil under essentially incompressible free-stream conditions as specified by the NASA Langley Research Center Turbulence Modeling Resource (<https://turbmodels.larc.nasa.gov>). The S-A turbulence model is used for this test case, and the free-stream conditions are as follows: Mach number of 0.09, angle of attack of 13.87° , Reynolds number based on a chord of 1.0 of $1.52e6$, and a free-stream temperature of 297.778 K. Mach contours around the NACA 4412 are shown in Fig. 5(a), and the small separation region on the upper trailing edge that is characteristic of this test case is visible. For additional SU2 validation results for the NACA 4412 case, the reader is again referred to previous work.¹⁵

The grid is composed entirely of quadrilateral elements in an O-grid topology with 60,996 total grid points and 60,398 cells. More specifically, it is a 598×102 curvilinear grid with 598 grid points along the airfoil surface and 102 grid points stretching from the airfoil surface to far-field boundary. The far-field boundary is located approximately 20 chord lengths away from the airfoil. A characteristic-based far-field condition is applied to the outer domain boundary, and an adiabatic, no-slip wall boundary condition is enforced on the airfoil surface. Mesh spacing at the airfoil boundary is sufficient to ensure $y^+ \approx 1$ over the airfoil surface.





Figure 6. NACA 4412 drag sensitivity with respect to vertical (y-axis) movement of 30 FFD control point variables.

We will again compare drag sensitivities with respect to vertical (y -axis) movement of FFD control point variables. Both the mesh and FFD control points are shown Fig. 5(b). Note that only the interior set of FFD control points are active, and the control points making up the outer boundary of the FFD box are fixed. The same set of numerical methods as with the RAE 2822 case were applied with one notable exception: no limiter is applied to the second-order upwind scheme in the primal problem.

Fig. 6 compares the drag sensitivities for the various adjoint formulations. SI FT shows decent agreement with AD and FD, while the signs for several control points are incorrect for both SI FT and FI FT. When it comes to comparing the formulations, the results are less conclusive here than with the RAE 2822 case. FI FT again corrects inaccuracies near the trailing edge of the airfoil, but FI FT also predicts different behavior at the leading edge compared to AD and FD. AD FT matches with AD and FD near the trailing edges, but it exhibits smaller values of sensitivity along the upper/lower surfaces (some errors in signs) and different behavior near the leading edge. However, note that the magnitudes of the gradients are much smaller overall relative to the RAE 2822 case. Without a shock, the drag sensitivities represent the impact of shape changes on the skin friction and pressure drag.

As a follow-up test, we compute the drag sensitivities again while varying the numerical methods. In practice, we find that the ability to choose and customize numerical methods for solving the continuous adjoint is typically a luxury, since convergence problems often arise with any type of adjoint. In these scenarios, we are able to tailor our numerical methods independently for both the primal and adjoint problems in order to obtain a viable solution for the continuous adjoint, albeit at the potential cost of accuracy or convergence. This is evident from Figs. 7(a)–7(c).



(a) Drag sensitivity with original set of numerical methods, i.e., second-order upwind without slope limiting. (b) Drag sensitivity after activating primal limiter with original adjoint numerics. (c) Drag sensitivity after switching to a centered scheme for the adjoint with original primal numerics.

Figure 7. NACA 4412 drag sensitivity results with changes to the numerical methods.

In Fig. 7(b), we activate the Venkatakrishnan limiter in the primal problem while maintaining the same numerical methods for the continuous adjoint problem. The limiting of the flow variable gradients causes a discrepancy in the SI FT result near the trailing edge. This might be expected, as sharp trailing edges do not have well-defined normals and are typically highly sensitive areas with large flow variable gradients, and both surface normals and gradients of variables on the surface are needed to evaluate Eqn. 1. For this case, the FI FT result recovers the behavior of AD and FD near the trailing edge.

In Fig. 7(c), we hold fixed the numerical methods for the primal problem but recompute the continuous adjoint with a second-order centered scheme in place of the second-order upwind scheme. Discrepancies appear at both the leading and trailing edges of the airfoil, and SI FT can be up to an order of magnitude away from the other gradient results. Again, it is not surprising that changes to the numerics can have a large impact on the variables appearing in Eqn. 1, as well as their spatial gradients near the surface, especially for RANS cases featuring highly-stretched, anisotropic cells for resolving boundary layers. The FI FT result for this case again recovers behavior similar to the AD and FD results.

Viewing the NACA 4412 results now as a whole, we see that our original set of numerical methods was a good choice for SI FT based on the general agreement seen in Fig. 6. In our experience, the SI FT formulation works well in most cases, especially when the grid is of high quality. Unfortunately, if we stray from this set of numerical methods, we can quickly find the SI FT result to be an order of magnitude different from the others. In these cases, we see that FI FT alleviates the issues, returning the continuous adjoint result in line with AD and FD. In that sense, the FI formulation offers additional robustness for the continuous adjoint, even when tailoring the numerical methods to a given problem, which is a common strategy.

D. Ahmed Body

The previous cases were simple and somewhat idealized, as the problems were 2D, the flow conditions were relatively tame, and the meshes were of high quality. Convergence to machine precision was easily obtained for all prior cases for both the primal and adjoint problems. Unfortunately, these luxuries can not always be expected in an industrial environment where turnaround times must fit within an established design cycle. Complex, 3D geometries embedded in coarse or low-quality unstructured meshes is often the norm, and the desired flow conditions may stress the intended limits of the available models.

With this in mind, we present a final case that is more representative of an industrial situation: turbulent flow around the 3D Ahmed car body configuration.³² The Ahmed body is a widely studied, canonical case for automotive applications, and while the geometry is not overwhelmingly complex, the flow conditions around the sharp-edged, bluff body result in a massively separated wake region. The S-A turbulence model is again used for this test case, and the free-stream conditions are as follows: Mach number of 0.1728, angle of attack of 0.0° , angle of sideslip of 0.0° , Reynolds number based on a body length of 1.044 m of $4.29e6$, and a free-stream temperature of 300.0 K. Pressure contours on the body surface and streamlines are shown in Fig. 8(a), and the separated wake region behind the body is visible.

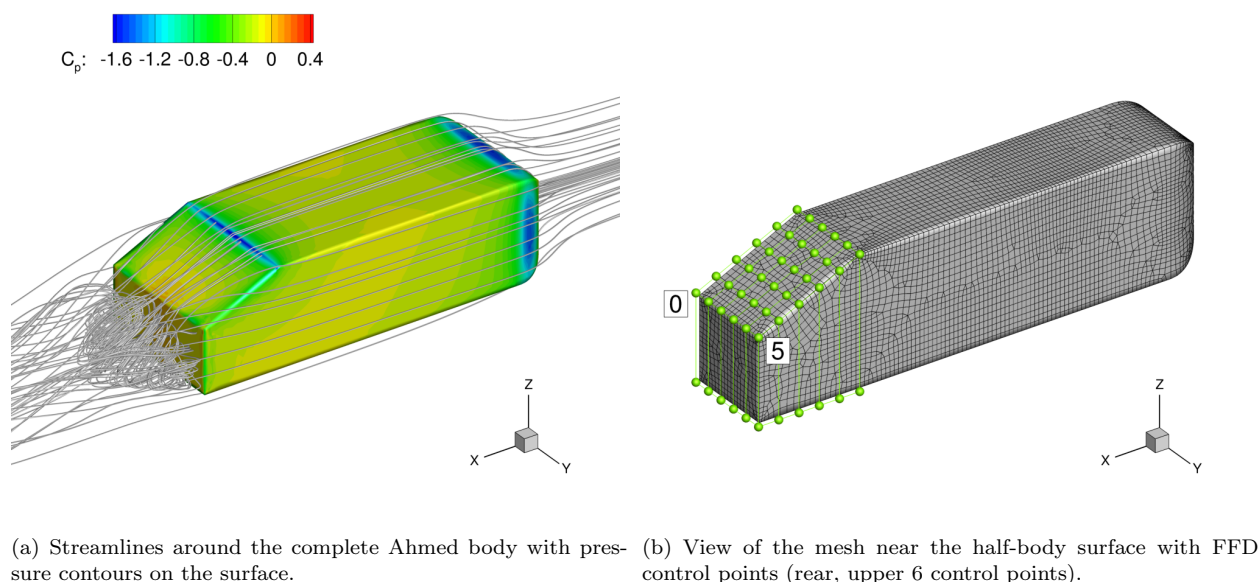


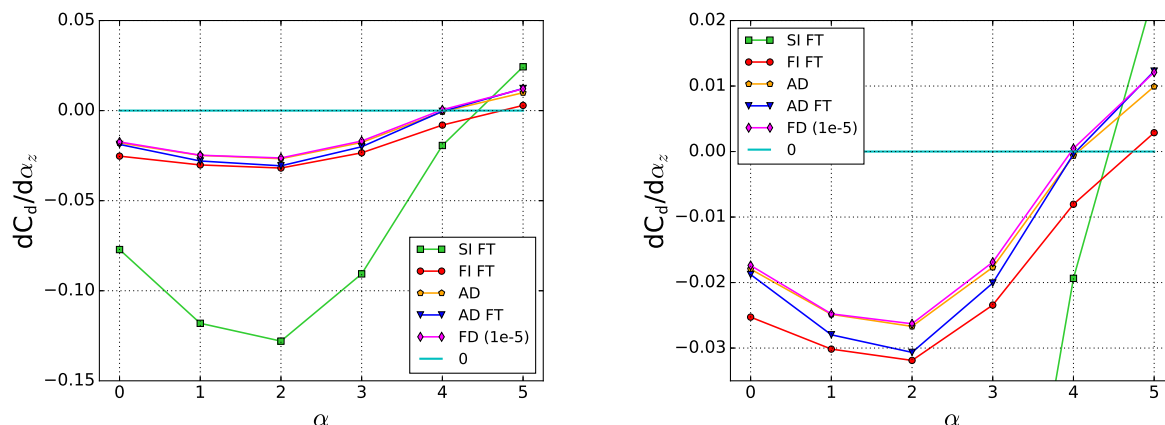
Figure 8. Streamlines around the Ahmed body geometry and view of the surface grid with FFD design variables.

For this case, we take advantage of the flow symmetry and create an unstructured grid for the half-body alone. The mixed-element grid has a combination of triangles and quadrilaterals on the body surface, a near-body region extruded normally from the body that is composed of prisms and hexas, and a far-field region of tets. The mesh contains a total of 247,862 grid points and 508,080 elements. The mesh is coarse for this problem by construction, especially in the wake region where extra mesh density is required to resolve the wake adequately. However, mesh spacing near the wall is sufficient to ensure that $y^+ \approx 1$.

The far-field boundary is located approximately 10 body lengths away, and a symmetry plane representing the ground is below the body. A characteristic-based far-field condition is applied to the outer domain boundary, a symmetry boundary condition is applied to the floor and the symmetry plane, and an adiabatic, no-slip wall boundary condition is enforced on the body. We will study the drag sensitivity with respect to 6 FFD control point variables along the upper, rear corner of the vehicle that are allowed to move in the vertical direction (z -axis). The surface mesh and FFD control points are shown Fig. 8(b).

We apply the JST scheme to the convective terms for the primal problem. The turbulent variable for the S-A model is convected with a first-order scalar upwind method, and viscous fluxes are calculated using the corrected average-gradient method. Implicit, local time-stepping is used to converge the problem to the steady-state solution, and the linear system is solved using GMRES. For the continuous adjoint problem, we use the modified JST scheme that features only a high-order dissipation term. It should be noted that we

were able to achieve convergence to machine precision for the primal problem, but the continuous adjoint was more difficult to converge with only 3 or 4 orders of magnitude reduction in the adjoint density residual.



(a) Drag sensitivity with respect to vertical (z-axis) movement of 6 FFD control point variables.

(b) Zoom view of the drag sensitivity.

Figure 9. Ahmed body drag sensitivity results.

Fig. 9(a) and 9(b) show the drag sensitivity comparisons in a full view and a zoom view, respectively. For this case, the results are conclusive: even though the signs are correct, SI FT can be an order of magnitude away from AD and FD, while FI FT more closely tracks AD and FD. We also see that the AD FT result moves from AD toward FI FT, which is intuitive if one considers that both AD FT and FI FT include the effects of mesh differentiation while freezing turbulence.

The Ahmed case features a coarse, mixed-element mesh with the necessary high aspect ratio cells for RANS, a large separated wake region, and lack of convergence for the continuous adjoint problem, but even in the face of these issues, FI FT provides robustness over SI FT and results in sensitivities that are generally close to AD and FD. Comparing the effects of freezing turbulence (AD vs. AD FT) and the inclusion of grid effects (SI FT vs. FI FT), we conclude that the inclusion of grid effects is more critical than the inclusion of an adjoint turbulence model for obtaining accurate sensitivities with this type of problem.

IV. Conclusions

In this study, we have investigated different adjoint formulations within the SU2 suite: a continuous adjoint in surface integral form with frozen turbulence (SI FT), a continuous adjoint in field integral form with frozen turbulence (FI FT), and a discrete adjoint using algorithmic differentiation with and without frozen turbulence (AD and AD FT). We have also compared sensitivity results against simple finite differences (FD). These formulations were applied to the computation of drag sensitivities with respect to shape changes for different external flow test cases. The four test cases considered were the NACA 0012 in transonic, inviscid flow, the RAE 2822 in transonic, turbulent flow, the NACA 4412 in subsonic, turbulent flow, and the Ahmed body in subsonic, turbulent flow. The final three cases all require the RANS equations, and two of them feature some level of separation.

From the cases studied, we are able to draw a number of conclusions about how the formulations perform. First, similar to the conclusions reached by Lozano et al.¹⁰ and Kavvadias et al.,¹³ we note that SI FT sensitivities can be more than an order of magnitude away from AD and FD in cases where the grid or numerical methods have a major impact on the solution. The FI formulation is much closer to AD and FD at sharp edges as compared to SI, which was discussed previously in Lozano,¹² but we also find that FI accuracy is more robust to changes in numerical methods over the SI formulation. Overall, FI FT is generally close to AD and FD for all cases studied, with occasional erroneous sensitivity signs.

Again, we would like to highlight that all of the above formulations are valid approaches to solving the adjoint problem and have been used successfully for aerodynamic shape design. Depending on the situation, any one of the above formulations may still be the best choice given constraints on time, memory, accuracy,

etc. However, from these results, we see that terms related to the mesh differentiation can add robustness to the continuous adjoint, and they can be more important than the inclusion of an adjoint turbulence model. In particular, grid effects appear to be more important in RANS cases than in inviscid cases, which is likely due to high aspect ratio cells near the walls for turbulent flows.

V. Acknowledgements

The first author would like to thank Carsten Othmer, Vaggelis Papoutsis-Kiachagias, and Kyriakos Giannakoglou for collaborating on earlier work that led to this paper and for their leading research in these topics. Antony Jameson deserves thanks for enlightening discussions and for sharing code related to the original FI formulation that guided the implementation in this paper. John Vassberg and Francisco Palacios are acknowledged for providing the NACA 0012 grids used in this study and for helpful advice along the way. Asitav Mishra and Karthik Duraisamy are thanked for collaborating on the Ahmed car body case.

References

- ¹O. Pironneau. On optimum design in fluid mechanics. *Journal of Fluid Mechanics*, 64:97–110, 1974.
- ²A. Jameson. Aerodynamic design via control theory. *Journal of Scientific Computing*, 3:233–260, 1988.
- ³S. Kim, J. J. Alonso, and A. Jameson. A gradient accuracy study for the adjoint-based Navier-Stokes design method. *AIAA Paper 1999-0299*, 1999.
- ⁴M. B. Giles and N. A. Pierce. An introduction to the adjoint approach to design. *Flow, Turbulence and Combustion*, 65:393–415, 2000.
- ⁵S. K. Nadarajah and A. Jameson. A comparison of the continuous and discrete adjoint approach to automatic aerodynamic optimization. *AIAA Paper 2000-0667*, 2000.
- ⁶K. Leoviriyakit, S. Kim, and A. Jameson. Viscous aerodynamic shape optimization of wings including planform variables. *AIAA Paper 2003-3498*, 2003.
- ⁷A. Jameson and S. Kim. Reduction of the adjoint gradient formula for aerodynamic shape optimization problems. *AIAA Journal*, 41(11):2114–2129, 2003.
- ⁸C. Castro, C. Lozano, F. Palacios, and E. Zuazua. A systematic continuous adjoint approach to viscous aerodynamic design on unstructured grids. *AIAA Journal*, 45(9):2125–2139, 2007. DOI: 10.2514/1.24859.
- ⁹D.I. Papadimitriou and K.C. Giannakoglou. A continuous adjoint method with objective function derivatives based on boundary integrals, for inviscid and viscous flows. *Computers & Fluids*, 36:325–341, 2007.
- ¹⁰C. Lozano, E. Andres, M. Martin, and P. Bitrian. Domain versus boundary computation of flow sensitivities with the continuous adjoint method for aerodynamic shape optimization problems. *Int. J. Numer. Meth. Fluids*, 70:1305–1323, 2012.
- ¹¹W. K. Anderson and V. Venkatakrishnan. Aerodynamic design optimization on unstructured grids with a continuous adjoint formulation. *AIAA Paper 1997-0643*, 1997.
- ¹²C. Lozano. Discrete surprises in the computation of sensitivities from boundary integrals in the continuous adjoint approach to inviscid aerodynamic shape optimization. *Computers & Fluids*, 56:118–127, 2012.
- ¹³I.S. Kavvadias, E.M. Papoutsis-Kiachagias, and K.C. Giannakoglou. On the proper treatment of grid sensitivities in continuous adjoint methods for shape optimization. *Journal of Computational Physics*, 301:1–18, 2015.
- ¹⁴F. Palacios, M. R. Colonno, A. C. Aranake, A. Campos, S. R. Copeland, T. D. Economon, A. K. Lonkar, T. W. Lukaczyk, T. W. R. Taylor, and J. J. Alonso. Stanford University Unstructured (SU²): An open-source integrated computational environment for multi-physics simulation and design. *AIAA Paper 2013-0287*, 2013.
- ¹⁵F. Palacios, T. D. Economon, A. C. Aranake, S. R. Copeland, A. K. Lonkar, T. W. Lukaczyk, D. E. Manosalvas, K. R. Naik, A. S. Padron, B. Tracey, A. Variyar, and J. J. Alonso. Stanford University Unstructured (SU²): Open-source analysis and design technology for turbulent flows. *AIAA Paper 2014-0243*, 2014.
- ¹⁶T. D. Economon, F. Palacios, S. R. Copeland, T. W. Lukaczyk, and J. J. Alonso. SU2: An open-source suite for multiphysics simulation and design. *AIAA Journal*, 54(3):828–846, 2016.
- ¹⁷R. M. Hicks and P. A. Henne. Wing design by numerical optimization. *Journal of Aircraft*, 15(7):407–412, 1978.
- ¹⁸J. A. Samareh. A survey of shape parameterization techniques. Technical Report NASA/CP-1999-209136, NASA, 1999.
- ¹⁹R. P. Dwight. Robust mesh deformation using the linear elasticity equations. Proceedings of the Fourth International Conference on Computational Fluid Dynamics (ICCFD), pages 401–406, 2006.
- ²⁰E. Jones, T. Oliphant, P. Peterson, et al. SciPy: Open source scientific tools for Python, 2001.
- ²¹A. Bueno-Orovio, C. Castro, F. Palacios, and E. Zuazua. Continuous adjoint approach for the Spalart-Allmaras model in aerodynamic optimization. *AIAA Journal*, 50(3):631–646, 2012.
- ²²F. Palacios, T. D. Economon, and J. J. Alonso. Large-scale aircraft design using SU2. *AIAA Paper 2015-1946*, 2015.
- ²³T. D. Economon, F. Palacios, and J. J. Alonso. Unsteady continuous adjoint approach for aerodynamic design on dynamic meshes. *AIAA Journal*, 53(9):2437–2453, 2015.
- ²⁴T. Albring, M. Sagebaum, and N.R. Gauger. Efficient aerodynamic design using the discrete adjoint method in SU2. *AIAA Paper 2016-3518*, 2016.
- ²⁵T. Albring, M. Sagebaum, and N. R. Gauger. Development of a consistent discrete adjoint solver in an evolving aerodynamic design framework. *AIAA Paper 2015-3240*, 2015.

- ²⁶A. Jameson, W. Schmidt, and E. Turkel. Numerical solution of the Euler equations by finite volume methods using Runge-Kutta time stepping schemes. *AIAA Paper 1981-1259*, 1981.
- ²⁷Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 7(3):856–869, 1986.
- ²⁸P. Cook, M. McDonald, and M. Firmin. Aerofoil RAE2822 pressure distributions, and boundary layer and wake measurements. Technical Report AGARD 138, 1979.
- ²⁹J. W. Slater, J. C. Dudek, and K. E. Tatum. The NPARC Alliance Verification and Validation Archive. *NASA Technical Report 2000-209946*, 2000.
- ³⁰P. L. Roe. Approximate riemann solvers, parameter vectors, and difference schemes. *Journal of Computational Physics*, 43(2):357–372, 1981.
- ³¹V. Venkatakrishnan. On the accuracy of limiters and convergence to steady state solutions. *AIAA Paper 1993-0880*, 1993.
- ³²Ahmed S. R., RAM G., and Faltin G. Some salient features of the time-averaged ground vehicle wake. *SAE Paper 840300*, 1984.