Please cite this document as:
Yu, Yin, Lyu, Zhoujie, Xu, Zelu, and Martins, Joaquim R. R. A., "On the Influence of Optimization Algorithm and Starting Design on Wing Aerodynamic Shape Optimizationl", *Aerospace Science and Technology*, 2018. doi: 10.1016/j.ast.2018.01.016.
This document can be found at: `http://mdolab.engin.umich.edu`.

On the Influence of Optimization Algorithm and Initial Design on Wing Aerodynamic Shape Optimization

Yin Yu[1], Zhoujie Lyu[2], Zelu Xu[3], Joaquim R. R. A. Martins[4]

Abstract Aerodynamic shape optimization is a useful tool in wing design, but the impact of the choice of optimization algorithm and the multimodality of the design space in wing design optimization is still poorly understood. To address this, we benchmark both gradient-based and gradient-free optimization algorithms for computational fluid dynamics based aerodynamic shape optimization problems based on the Common Research Model wing geometry. The aerodynamic model solves the Reynolds-averaged Navier–Stokes equations with a Spalart–Allmaras turbulence model. The drag coefficient is minimized subject to lift, pitching moment, and geometry constraints, with up to 720 shape variables and 11 twist variables for two mesh sizes. We benchmark six gradient-based and three gradient-free algorithms by comparing both the accuracy of the optima and the computational cost. Most of the optimizers reach similar optima, but the gradient-based methods converge to more accurate solutions at a much lower computational cost. Since multimodality and nonsmoothness of the design space are common arguments for the use of gradient-free methods, we investigate these issues by solving the same optimization problem starting from a series of randomly generated initial geometries, as well as a wing based on the NACA 0012 airfoil with zero twist and constant thickness-to-chord ratio. All the optimizations consistently converge to practically identical results, where the differences in drag are within 0.05%, and the shapes and pressure distributions are very similar. Our overall conclusion is that the design space for wing design optimization with a fixed planform is largely convex, with a very small flat region that is multimodal because of numerical errors. However, this region is so small, and the differences in drag so minor, that the design space can be considered unimodal for all practical purposes.

1 Introduction

The aerodynamic shape optimization of transonic aircraft wings has long been a difficult and expensive task. Small changes in shape can have a large impact on aerodynamic performance, and therefore the optimization requires hundreds of design variables Lyu and Martins (2014). Thus, aerodynamic shape optimization based on computational fluid dynamics (CFD) can be costly.

Aerodynamic shape optimization problems can be solved with gradient-based or gradient-free methods. Gradient-based methods are preferable when an efficient gradient evaluation is available Jameson (2003). The application of gradient-based optimization to this problem was pio-

[1]Aerospace Engineer PSA Branch, 21000 Brookpark Rd., MS 5-11; Doctoral Candidate; Department of Aerospace Engineering, University of Michigan; AIAA Member
[2]Research Investigator, Department of Aerospace Engineering, AIAA Senior Member
[3]Aerospace Engineer, Science and Technology Corporation; AIAA Senior Member
[4]Professor, Department of Aerospace Engineering; AIAA Associate Fellow

neered in the 1970s, with gradients computed using finite-difference approximations Hicks and Henne (1978). As the number of design variables increases, the cost of this computation becomes prohibitive. Adjoint methods were developed to address this issue; they provide a way to evaluate the gradients with a cost that is independent of the number of design variables. Peter and Dwight Peter and Dwight (2010) reviewed these and other methods for computing aerodynamic shape derivatives. Martins and Hwang Martins and Hwang (2013) generalized the adjoint method and discussed its connection to other derivative evaluation methods.

Pironneau pioneered the use of adjoint-based gradient calculation in airfoil profile optimization by deriving the adjoint for the Stokes equations Pironneau (1973) and for the incompressible Euler equations Pironneau (1974). Jameson (1988) then made the adjoint method useful in the design of transonic airfoils by developing an adjoint for inviscid compressible flow. The aerodynamic design of transonic wings requires a model that can represent the shock-wave boundary layer interaction, since there is a strong nonlinear coupling between airfoil shape, wave drag, and viscous effects. Therefore, transonic wing optimization based on the Euler equation performs poorly when analyzed in turbulent flow Lyu et al. (2013, 2015).

The adjoint method was later extended to the compressible Navier–Stokes equations with turbulence models, making it possible to solve practical aerodynamic design problems. Jameson et al. Jameson et al. (1998) optimized a wing-body configuration modeled with the compressible Navier–Stokes equations using a continuous adjoint approach. They used a 590k-cell mesh and achieved a shock-free solution at Mach 0.86. Anderson and Venkatakrishnan (1999) optimized airfoil shapes using a discrete adjoint that included the linearization of the Spalart–Allmaras turbulence model. Nielsen and Anderson (1999) further extended the approach to the three-dimensional Reynolds-averaged Navier–Stokes (RANS) equations. They achieved an 8% drag reduction for the ONERA M6 wing with thickness and camber design variables at two chordwise locations. Dwight and Brezillon (2009) and Brezillon and Dwight (2012) optimized the DLR-F6 wing-body configuration using a RANS solver and a discrete adjoint, achieving a 10-count drag reduction by varying 96 design variables.

Lyu et al. (2013) developed a discrete adjoint for the RANS equations and Spalart–Allmaras turbulence model using automatic differentiation to construct the required derivative terms. They used this adjoint implementation to perform aerodynamic shape optimizations of the ONERA M6 wing with 192 design variables for both the Euler and RANS models. They observed significant differences between the optimal shapes obtained with Euler and RANS, which emphasized the importance of including the viscous compressible effects in transonic aerodynamic shape design. The framework developed by Lyu et al. (2013) has since been used in a variety of applications and studies Kenway and Martins (2017); Liem et al. (2017); Chen et al. (2016); Kenway and Martins (2016, 2015). Telidetzki et al. (2014) performed a series of high-fidelity aerodynamic shape optimizations using a parallel Newton–Krylov–Schur method based on the Euler or RANS equations. They demonstrated the effectiveness of the gradient-based aerodynamic shape optimization methodology, obtaining significant drag reductions in all their cases. Chen et al. (2016) performed RANS-based aerodynamic shape optimization on a common research model (CRM) wing-body-tail configuration. Elham (2015) presented a quasi-three-dimensional method for wing aerodynamic analysis and drag prediction. They used a combination of the adjoint method, the chain rule for differentiation, and automatic differentiation to compute the gradients. Drela (1998, 1993) performed a constrained shape optimization on two-dimensional airfoils, using the Newton-based direct method to generate sensitivity information from inviscid Euler equations.

The gradient-free methods are generally easier to implement and use, and several of them are geared toward finding global optima. However, they incur a higher computational cost compared with gradient-based methods, especially when costly high-fidelity simulations are involved. Genetic

algorithms (GA) and their derivatives are among the most widely used gradient-free methods today Jones and Finch (1984); Marco et al. (1997). GAs are particularly suitable for problems with discontinuous objective functions, discrete design variables, or multiple local optima, i.e., multimodal functions. He and Agarwal (2014) performed aerodynamic shape optimization of a wind turbine blade airfoil using a multiobjective GA.

There have been a few studies of the performance of different optimizers for aerodynamic shape optimization. Zingg et al. (2008) compared gradient-based methods and a GA in aerodynamic airfoil optimization. They found that the GA used 5 to 200 times more function evaluations than the gradient-based method to find the optimum design. They suggested that GAs are better suited for low-fidelity preliminary design, while gradient-based methods are preferable for high-fidelity detailed design. Obayashi and Tsukahara (1997) compared a gradient-based method with simulated annealing and a GA on an airfoil lift maximization problem. The GA required the highest number of function evaluations but achieved the best design.

Gradient-based methods can converge to a local minimum when the objective or constraint functions involved are multimodal. Holst and Pulliam (2001) and Sasaki et al. (2001) both used GAs for airfoil and wing optimization cases, and they found no evidence of multimodality. Chernukhin and Zingg (2013) compared the performance of a gradient-based method, a GA, and a hybrid approach on a two-dimensional airfoil shape optimization and three-dimensional wing optimizations based on the Euler equations. While they concluded that the airfoil design problem was unimodal, they found multiple local optima for the wing case. In addition to twist and airfoil shape variables, the wing optimization cases included planform variables (chord variation, sweep, and dihedral). The physical significance of these multiple local optima is compromised by the fact that no viscous effects were considered. Therefore, variations in surface area and local chord do not affect drag as they would in the real design problem, leading to a design space that is completely different from the true physical one. Furthermore, dihedral has a weak influence on the aerodynamic forces, and letting dihedral vary without a penalty on the viscous drag leads to designs that are not realistic. A more recent study by Bons et al. (2017) has started to address multimodality with respect to planform variables as well.

Lyu et al. (2015) solved the AIAA Aerodynamic Design Optimization Discussion Group (ADODG) CRM wing using a gradient-based RANS solver[1]. This problem involves a lift-constrained drag minimization, where the design variables are the spanwise twist distribution and airfoil shapes. They achieved a 8.5% drag reduction using a multilevel optimization approach, and they addressed multimodality concerns by starting the same optimization problem from randomly generated initial geometries. They observed multiple local optima around a small region, but these were close together and exhibited similar drag values. Other researchers have also tackled this problem. Dumont and Méheut (2016) analyzed the optimal geometries obtained by Lyu et al. (2015) with their solver and independently verified the performance of this design, adding further insight using their drag decomposition tool. Lee et al. (2015) obtained similar results and did not report multiple local minima for this problem. Shi-Dong et al. (2017) also solved the ADODG CRM wing and concluded that all the results point to a unimodal design space for the CRM wing. Finally, Koo and Zingg (2018) performed another study of the ADODG CRM case, and they concluded that it does not have multiple local optima.

Motivated by the work cited above, our goals are twofold: to compare various gradient-based and gradient-free optimizers, and to examine the issue of multiple local minima more closely. We focus on the ADODG CRM design optimization mentioned above, which does not include planform design variables Lyu et al. (2015). Once the planform is allowed to vary, many other issues arise, and it is difficult to obtain a meaningful design optimization problem without considering other aircraft design aspects, such as structural weight and stability. We benchmark several optimization

algorithms using a wing twist optimization problem and a wing shape problem. Six of the optimizers are gradient-based and three are gradient-free.

To examine the issue of multiple local minima, we perform various optimizations starting from several random initial points. We also use an initial geometry that has the planform of a CRM wing but with zero initial twist and a NACA 0012 airfoil. We go beyond the study of Lyu et al. (2015) by trying different variations in the design variable set. We also look more closely at the cluster of close local minima by using even smaller convergence tolerances and by performing a grid refinement.

2 Numerical tools

We now describe the numerical methods and tools that are used for this study. These tools are a subset of the multidisciplinary design optimization (MDO) framework of aircraft configurations with high fidelity (MACH) Kenway et al. (2014). MACH can perform the simultaneous optimization of aerodynamic shape and structural sizing variables considering aeroelastic deflections Kenway and Martins (2014). In the present work, we use only the components of MACH that are relevant for aerodynamic shape optimization: the geometric parametrization, mesh perturbation, CFD solver, and optimization algorithm.

2.1 Geometric parametrization

We use a free-form deformation (FFD) volume approach to parametrize the wing geometry Kenway et al. (2010). Although we did not study the effect of different parameterizations, Lee et al. (2017) did a study that compared FFD volumes and B-splines. They found that the two parameterizations converged to similar results, with drags within 0.2 counts of each other. Given such a small difference, they concluded that the design space is very flat around the optimum, which, as we will see, is consistent with our observations.

The FFD volume parametrizes the geometry changes rather than the geometry itself, resulting in a more efficient and compact set of geometry design variables, thus making it easier to manipulate complex geometries. We may embed any geometry inside the volume by performing a Newton search to map the parameter space to the physical space. All the geometric changes are performed on the outer boundary of the FFD volume. Any modification of this outer boundary indirectly modifies the embedded objects. The key assumption of the FFD approach is that the geometry has a constant topology throughout the optimization process, which is usually the case in wing design. In addition, since FFD volumes are trivariate B-spline volumes, the derivatives of any point inside the volume can be computed efficiently and accurately. Figure 1 shows the FFD volume and the geometric control points used in the aerodynamic shape optimization. The shape design variables are the displacement of all FFD control points in the vertical (z) direction.

2.2 Mesh perturbation

Since FFD volumes modify the geometry during the optimization, we must perturb the mesh for the CFD to solve for the modified geometry. The mesh perturbation scheme used in this work is a hybridization of the algebraic and the linear elasticity methods, developed by Kenway et al. (2010). The idea behind the hybrid scheme is to apply a linear-elasticity-based perturbation scheme to a coarse approximation of the mesh to account for large low-frequency perturbations, and to use the algebraic warping approach to attenuate small high frequency perturbations. For the results in this paper, the additional robustness of the hybrid scheme is not required, so we use only the algebraic scheme.

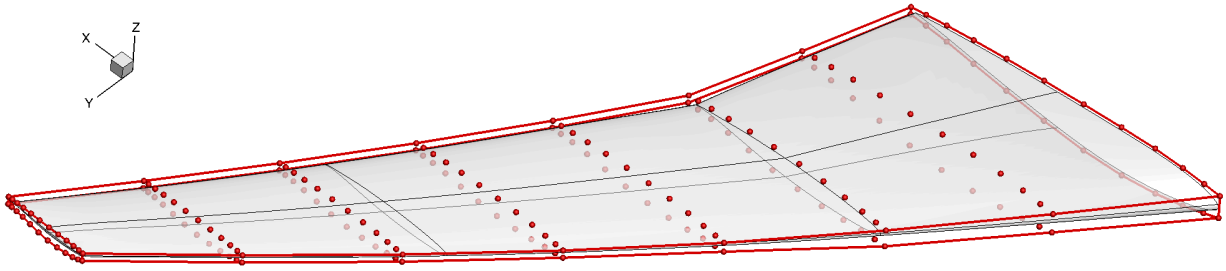

Figure 1: FFD volume showing the location of 192 control points (red spheres).

2.3 CFD solver

We use ADflow Lyu et al. (2013); van der Weide et al. (2006) as the CFD solver. It is a finite-volume, cell-centered multiblock solver for the compressible Euler, laminar Navier–Stokes, and RANS equations (steady, unsteady, and time-periodic). ADflow provides options for a variety of turbulence models with one, two, or four equations, and options for adaptive wall functions. The Jameson–Schmidt–Turkel (JST) scheme Jameson et al. (1981) augmented with artificial dissipation is used for the spatial discretization. The main flow is solved using an explicit multi-stage Runge–Kutta method, along with a geometric multigrid. A segregated Spalart–Allmaras turbulence equation is iterated with the diagonally dominant alternating direction implicit (DDADI) method.

To efficiently compute the gradients required for the optimization, we have developed and implemented a discrete adjoint method for the Euler and RANS equations within ADflow Mader et al. (2008); Lyu et al. (2013); Martins et al. (2016). The adjoint implementation supports both the full-turbulence and frozen-turbulence modes, but in the present work we use the full-turbulence adjoint exclusively. The adjoint is verified against the complex-step method Martins et al. (2003). We solve the adjoint equations with preconditioned GMRES Saad and Schultz (1986) using PETSc Balay et al. (1997, 2013a,b). We have previously performed extensive Euler-based aerodynamic shape optimization Mader and Martins (2013); Lyu and Martins (2014) and aerostructural optimization Kenway and Martins (2014); Liem et al. (2015). However, we have observed serious issues with Euler-based optimal designs: while Euler-based optimization can provide some design insight, the optimal Euler shapes are significantly different from those obtained with RANS Lyu et al. (2013). Euler-optimized shapes tend to exhibit a sharp pressure recovery near the trailing edge, which is nonphysical because such conditions near the trailing edge would cause separation. Thus, RANS-based shape optimization is necessary to achieve realistic designs.

2.4 Optimization algorithms

A number of optimizers are studied in this paper. We use the pyOpt framework Perez et al. (2012), which is an open-source framework that provides a common interface to all optimizers. In this section, we briefly describe each optimizer used.

SNOPT implements a sequential quadratic programming (SQP) method Gill et al. (2002). To solve the QP subproblems, it uses a smooth reduced-Hessian semidefinite QP solver with an augmented Lagrangian merit function. It solves large-scale design problems with a smooth objective and nonlinear constraints.

SLSQP is a sequential least squares programming algorithm Kraft (1988) that evolved from the

least squares solver of Lawson and Hanson Lawson and Hanson (1987). It uses a quasi-Newton Hessian approximation and an L1-test function in the line search algorithm. Kraft Kraft (1994) applied this method to aerodynamic and robotic trajectory optimization.

PSQP is a preconditioned SQP method with a Broyden–Fletcher–Goldfarb–Shanno (BFGS) variable metric update. It can handle large-scale problems with nonlinear constraints.

IPOPT implements a primal-dual interior-point algorithm Waechter and Biegler (2006) with a filter line search method Fletcher and Leyffer (2002). The barrier problem is solved by applying a damped Newton's method to the primal-dual equations. The line search method includes a second-order correction.

CONMIN solves linear or nonlinear constrained optimization problems using the method of feasible directions Vanderplaats (1978). It minimizes the objective function until it reaches an infeasible region. The optimization then continues by following the constraint boundaries in a descent direction.

GCMMA is a modified version of the method of moving asymptotes, designed for nonlinear programming and structural optimization Svanberg (1987). It solves a strictly convex approximating subproblem at each iteration. GCMMA guarantees convergence to a local minimum from any feasible initial point.

ALPSO is a parallel augmented Lagrange multiplier particle swarm optimization (PSO) solver implemented in Python Jansen and Perez (2011). PSO methods can solve nonsmooth objective functions and are more likely to find the global minimum. Augmented Lagrange multipliers are used to handle the constraints. ALPSO can be used for nonlinear, nondifferentiable, and nonconvex problems. It can also dynamically adjust the number of iterations in the inner steps to ensure the efficient solution of large-scale problems. Perez and Behdinan Perez and Behdinan (2007) applied this method to a nonconvex constrained structural problem. Other applications include the aeroservoelastic design optimization of a flexible wing Haghighat et al. (2012) and the aerostructural optimization of nonplanar lifting surfaces Jansen et al. (2010).

NSGA2 is a nondominated sorting-based multiobjective evolutionary algorithm Deb et al. (2002). The optimizer enforces constraints by tournament selection. It can solve nonsmooth and nonconvex multiobjective functions and tends to approach the global minimum.

NOMAD is a C++ implementation of the mesh adaptive direct search (MADS) algorithm Audet et al. (2006) for blackbox optimization with nonlinear constraints Le Digabel et al. (2012); Le Digabel (2011). It performs an adaptive search on a tower of underlying meshes to find a better solution and also controls the refinement of the meshes.

3 Problem formulation

Our goal is to perform lift-constrained drag minimization of the NASA CRM wing using the RANS equations as the aerodynamic model. In this section, we provide a detailed description of the problem.

3.1 Baseline geometry

The baseline geometry is a wing with a blunt trailing edge extracted from the CRM wing-body geometry. The NASA CRM geometry was developed for applied CFD validation studies Vassberg

Grid level	Grid size	$y+$
L2	450,560	2.213
L3	56,320	8.409

Table 1: Grid size used in aerodynamic twist optimization

et al. (2008). The CRM is representative of a contemporary transonic commercial transport, with a size similar to that of a Boeing 777. The CRM has 3.5 deg more quarter-chord wing sweep and 10.3% less wing area than the Boeing 777-200. The CRM geometry has been optimized for aerodynamic performance. However, several design features, such as an aggressive pressure recovery in the outboard wing, were introduced into the design to make it more interesting for research purposes and to protect intellectual property. This baseline geometry provides a reasonable initial point for the optimization, while leaving room for further performance improvements. In addition, the CRM was designed together with the fuselage of the full CRM configuration, so its performance is degraded when only the wing is considered. The geometry and specifications are given by the ADODG, and we repeat them here for convenience. The fuselage and tail are removed from the original CRM, and the root of the remaining wing is moved to the symmetry plane.

This baseline geometry is shown in Figure 2. All the coordinates are scaled by the mean aerodynamic chord (275.8 in). The resulting reference chord is 1.0, and the half span is 3.758151. The moment reference point is at $(x, y, z) = (1.2077, 0.0, 0.007669)$, while the reference area is 3.407014.

3.2 Mesh generation

We generate the mesh for the CRM wing using an in-house hyperbolic mesh generator. The mesh is marched out from the surface mesh using an O-grid topology to a farfield located at a distance of 25 times the span (about 185 mean chords). The nominal cruise flow condition is Mach 0.85 with a Reynolds number of 5 million based on the mean aerodynamic chord. The mesh we generated is an L0-level grid that contains 28.8 million cells. We then coarsen it to the L2 level (450 thousand cells) and the L3 level (56 thousand cells) for the optimization cases in this paper, as shown in Figure 3 and Figure 4. The meshes were generated using a hyperbolic mesh generator. The mesh size and $y^+_{\max}$ values under the nominal flight condition are listed in Table 3.2.

3.3 Optimization problem formulation

The aerodynamic shape optimization seeks to minimize the drag coefficient subject to a lift constraint ($C_L = 0.5$) and a pitching moment constraint ($C_{My} \geq -0.17$) at a cruise Mach of 0.85.

The shape changes are controlled by three sets of design variables: angle of attack, airfoil shape, and twist variables. The angle of attack is primarily used to match the lift constraint. A total of up to 192 shape variables are distributed on the lower and upper surfaces of the FFD volume, as shown in Figure 1. We found this number of shape variables to be the best trade-off between computational cost and accuracy in the optimal shape Lyu et al. (2015). We also use 10 sectional twist design variables that are crucial in tailoring the spanwise twist distribution. The center of the twist rotation is fixed at the reference axis, which is located at the quarter chord of each section. Figure 5 shows how the design variables are distributed.

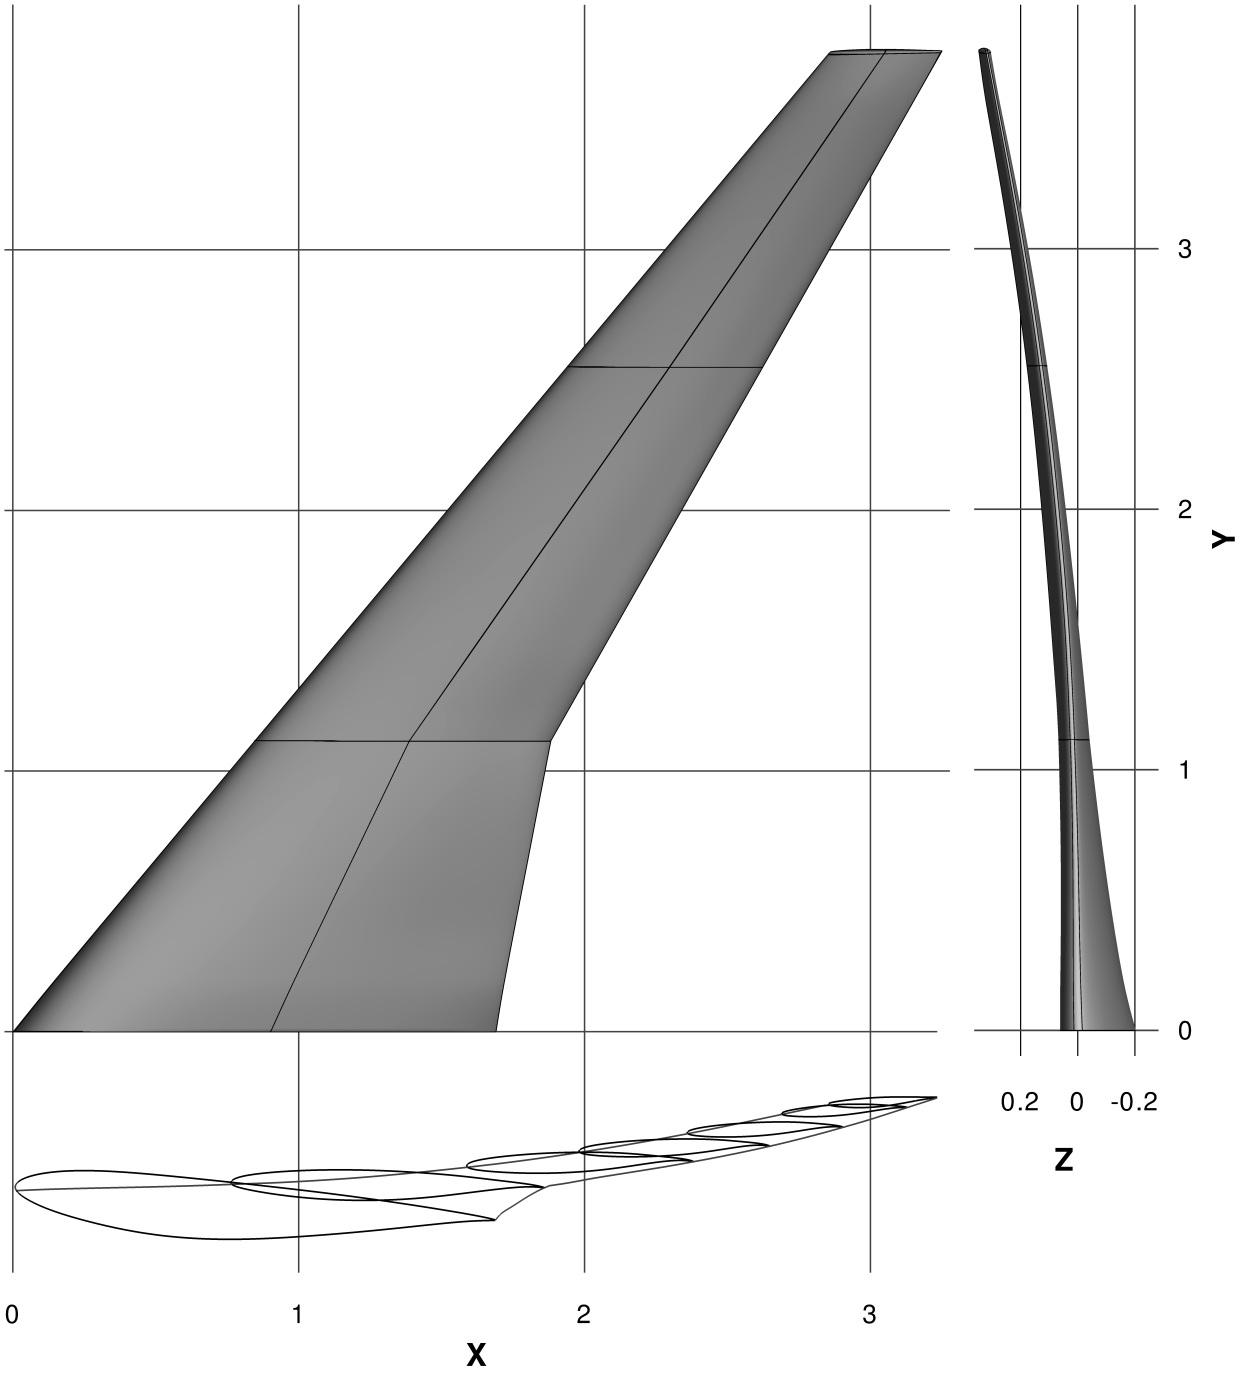

Figure 2: Common Research Model wing geometry.

There are 750 thickness constraints imposed in a 25 chordwise by 30 spanwise grid of points covering the full span and from 1% to 99.5% local chord. The control points at the trailing edge are constrained to avoid any movement of this edge. The leading-edge control points at the wing root are also constrained to maintain a constant incidence for the root section. The thickness is enforced to be greater than 25% of the baseline thickness at each location. While this value is unrealistically

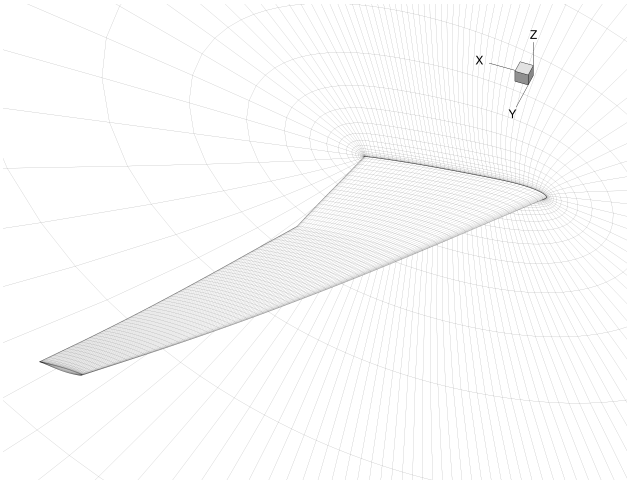

Figure 3: L2 mesh: 450 thousand cells

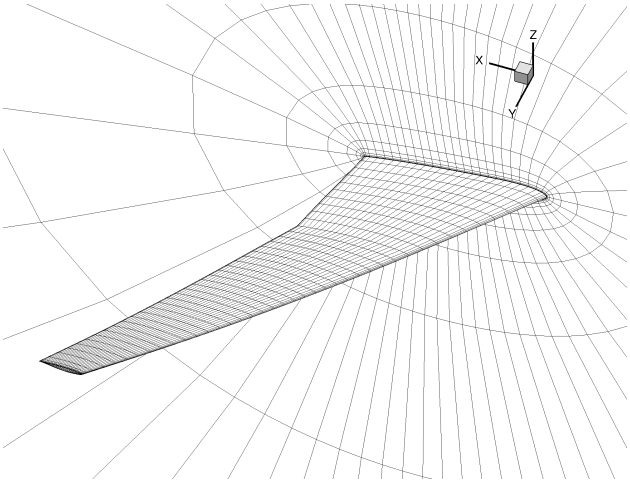

Figure 4: L3 mesh: 56 thousand cells

small (the wing would either fail structurally or be too heavy), letting the optimizer change the wing so much provides a better test of the optimization algorithms, shape parametrization, mesh movement, and derivative computations. The internal volume is constrained to be greater than or equal to the baseline volume. The pitching moment constraint is not applied to the cases in Section 4, because it is too costly for the gradient-free methods.

4 Optimization algorithm benchmarking

All the optimizations are performed at the nominal flight condition (Mach 0.85, $Re = 5 \times 10^6$), and the flow is solved using the RANS equations. We applied the lift coefficient constraint $C_L = 0.5$ and neglected the moment constraint in the original ADODG problem formulation to achieve a reasonable run time for the gradient-free optimizations.

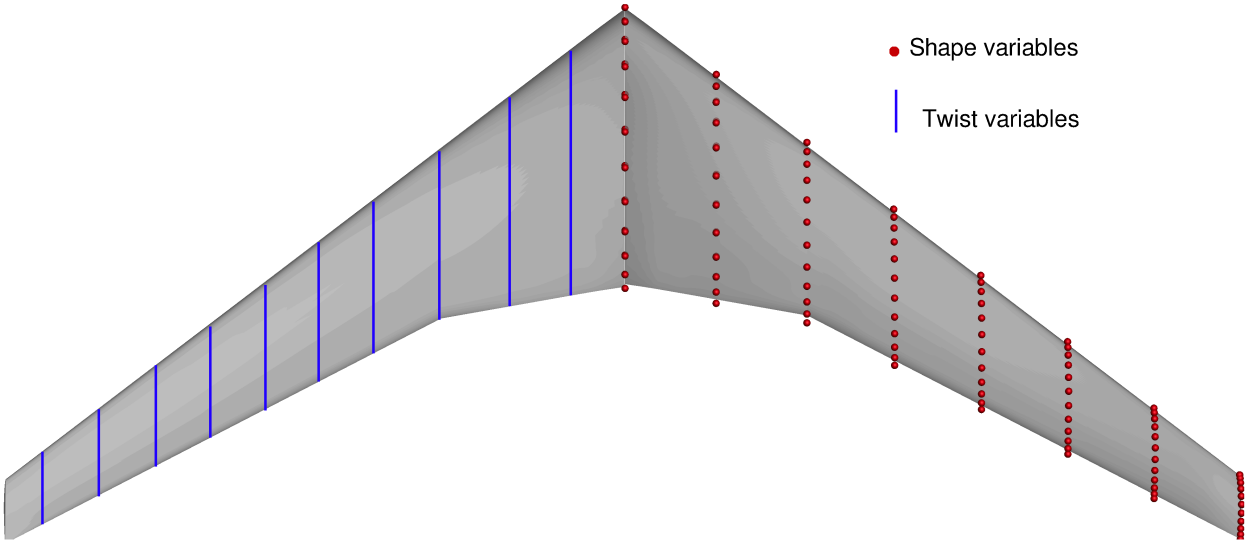

Figure 5: Shape and twist design variables.

4.1 Aerodynamic twist optimization

In this problem, we exclude the shape variables and consider only eight twist design variables evenly distributed on the wing, to limit the number of design variables so that we can run the gradient-free optimizations. The coarser L3 grid is used to reduce the optimization run time. In addition, we performed the optimization on the L2 grid for the gradient-based optimizers.

We compare the optimized twist, spanwise lift, and pressure distribution for each optimizer on the L3 grid in Figure 6. All the optimizers except NSGA2 converge to a similar drag value; the differences are within 0.1 drag counts (1 drag count $= 10^{-4} \times C_D$), and the corresponding twist distributions are nearly identical.

The computational costs for the optimizers are compared in Table 2, where they are ordered from fastest to slowest. The relative convergence tolerances for the ALPSO optimizer are 10^{-2} for the objective and 10^{-3} for the constraints, while for the gradient-based methods, these tolerances are 10^{-5} and 10^{-4}, respectively. The gradient-free optimizers need two to three orders of magnitude more iterations than the gradient-based optimizers. In this case, the gradient-based optimizers SNOPT, SLSQP, PSQP, and IPOPT perform well. GCMMA is slightly slower, and CONMIN is much slower. Of the three gradient-free optimizers, NOMAD performs the best, requiring 16 times the iterations of the best gradient-based optimizers, while NSGA2 performs the worst, requiring more than 500 times the computational effort. NOMAD and ALPSO converge to the same optimum as the gradient-based methods, but NSGA2 converges to a different twist distribution with a drag that is 0.82 counts higher.

Figures 7 and 8 show the convergence history of each optimizer. We plot the convergence of the gradient-based and gradient-free methods separately, since the numbers of iterations for the former are two orders of magnitude lower. For the gradient-based methods, we plot the value of the objective function with respect to the number of function evaluations. For the gradient-free

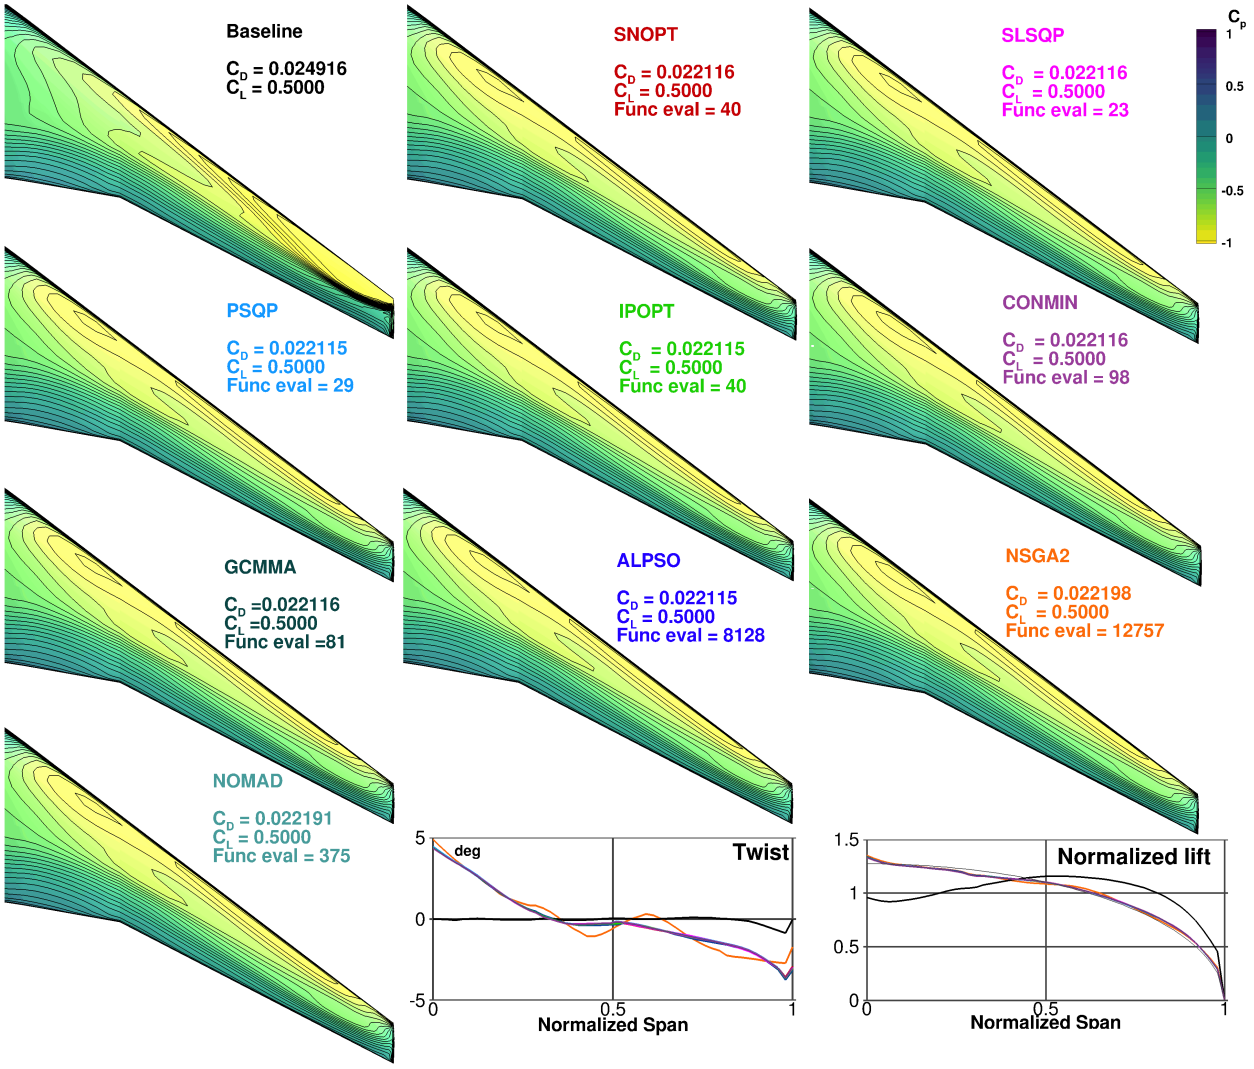

Figure 6: Aerodynamic twist optimization comparison on the L3 grid.

methods, we plot only the best point of each iteration or generation.

We verify the optimization using the finer L2 grid. We use only the gradient-based methods in this verification because of the prohibitive cost of the gradient-free methods. Figure 9 shows the optimized results on the L2 grid. The shock on the wing cannot be removed completely with only the twist design variables. The drag is reduced by 29 counts, and the difference in drag between the optimizers is within 0.1 counts. Therefore, it appears that the twist problem has only one optimum. The different grid resolution results in a similar optimized lift distribution but a different optimized twist distribution, as shown in Figure 10. The optimized design increases lift at the root and reduces lift at the tip, thus moving toward an elliptical lift distribution. However, since the optimizers minimize the total drag with only eight twist design variables, the optimal trade-off between induced drag with wave and viscous drag is not obvious, resulting in a nonelliptical lift distribution.

In this study, we examined a twist design problem. We used eight design variables subject to a

Optimizer	Drag (counts)	Iterations	Proc-hours
SLSQP	221.16	23	2.56
IPOPT	221.15	40	2.69
PSQP	221.15	29	3.70
GCCMA	221.16	81	4.57
SNOPT	221.16	40	5.81
CONMIN	221.16	98	33.61
NOMAD	221.91	375	40.23
ALPSO	221.15	8128	1695.72
NSGA2	221.98	12757	2744.16

Table 2: Computational cost comparison for the twist optimization problem using the L3 grid.

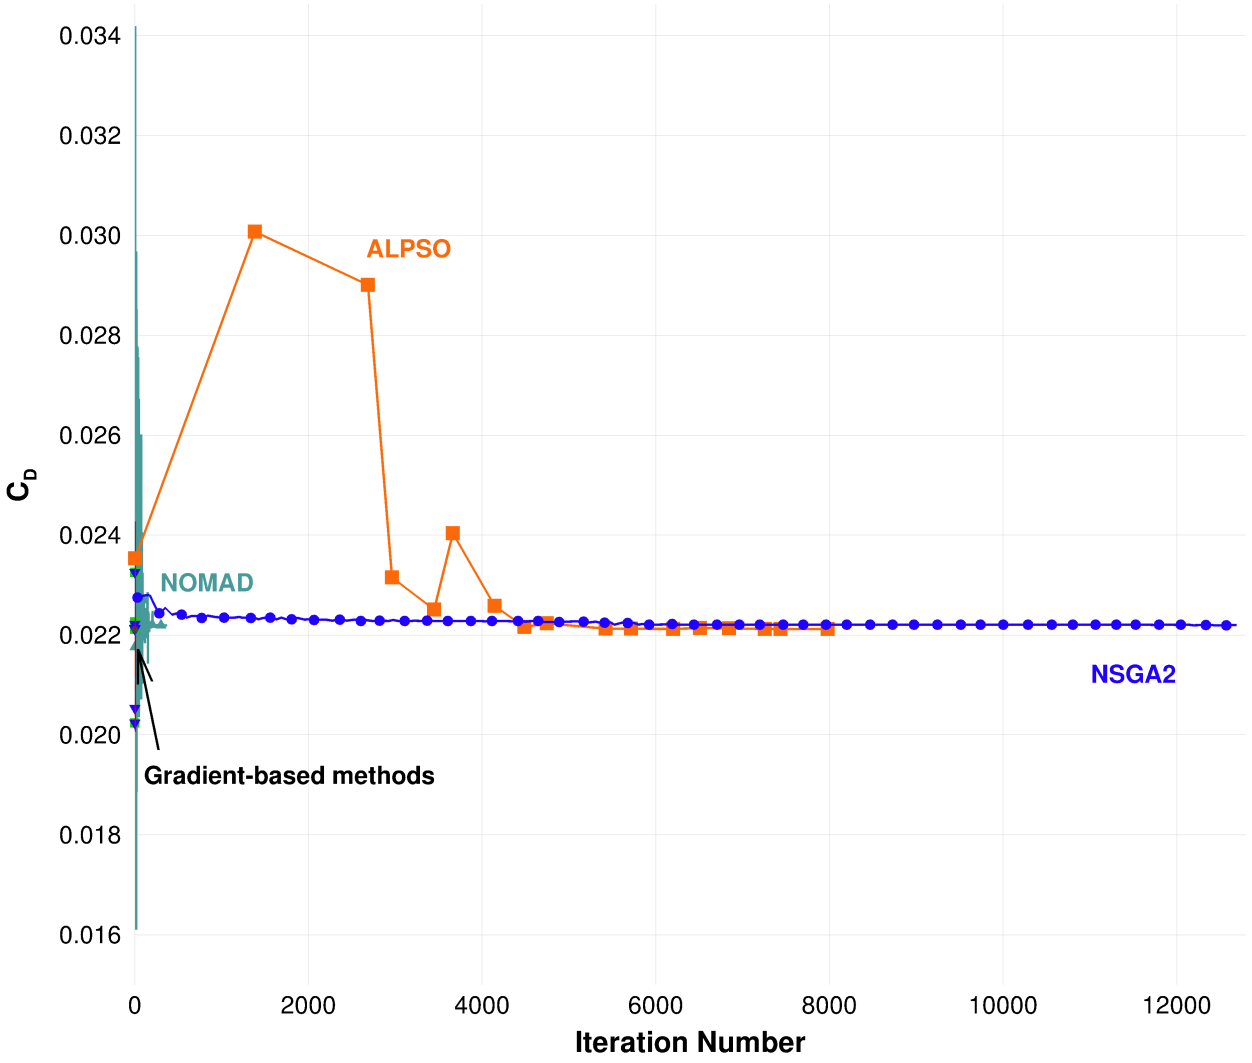

Figure 7: Twist optimization convergence history of the gradient-free methods for the L3 grid Lyu et al. (2014).

lift constraint. We compared the optimized results using different optimizers on two grid levels. All

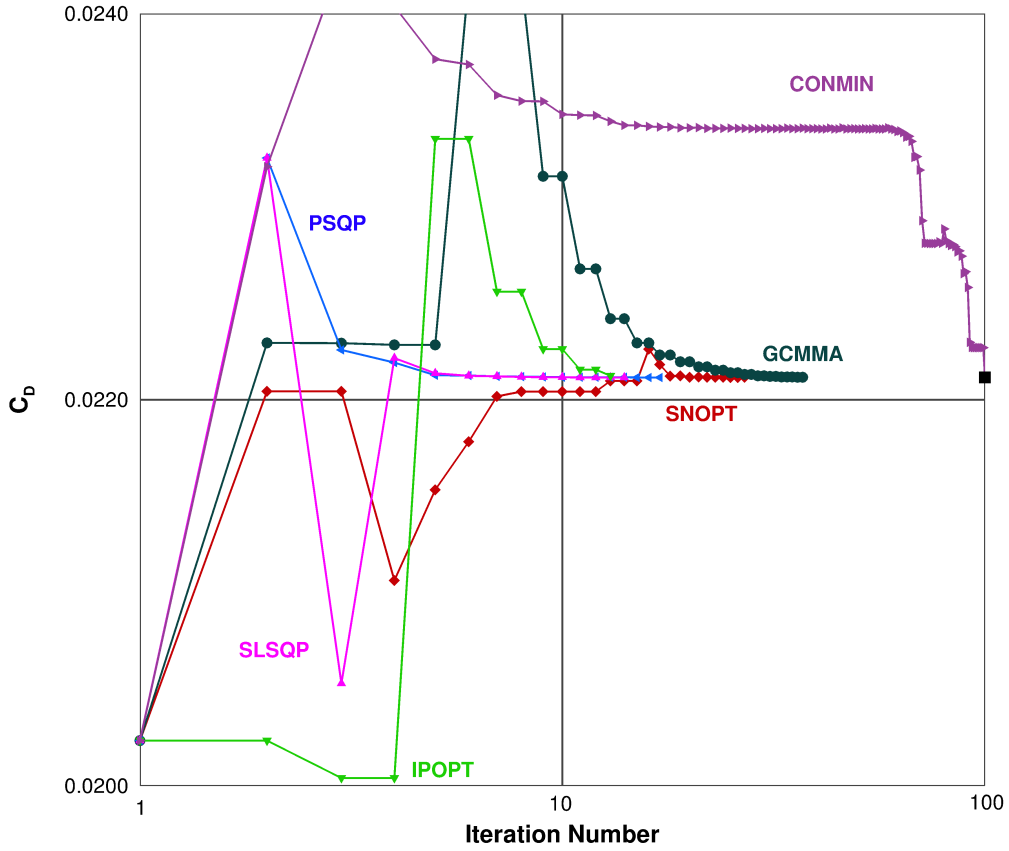

Figure 8: Twist optimization convergence history of the gradient-based methods for the L3 grid Lyu et al. (2014).

the optimizers converged to similar optima. A single global minimum is observed. The gradient-based methods converged significantly faster than the gradient-free methods, and thus they are better choices for high-fidelity aerodynamic shape optimization.

4.2 Aerodynamic shape optimization

In this study, we use the same baseline geometry discussed above. However, we use a total of 192 shape design variables instead of 8 twist design variables. The lift, thickness, and volume constraints are as stated in Section 3.3. This problem requires significantly more computational power than the previous case, so we use the L2 grid and only four of the gradient-based optimizers. The convergence tolerance is 10^{-6} for the objective and 10^{-4} for the constraints.

Figure 11 shows results from the four optimizers. The results from the baseline geometry are shown in black. The drag is reduced by 4.84%, from 206.7 to 196.6 counts.

All the optimizers are able to completely eliminate the shock on the wing and achieve an elliptical lift distribution. While this is not a desirable distribution in reality because of structural considerations, convergence to this theoretical result provides a good validation of the methods. For more practical design optimization, it is desirable to include structural stress constraints and to optimize both the wing shape and the structural sizing to obtain wing designs that represent optimal trade-offs between aerodynamic and structural performance Kennedy et al. (2014); Kenway and Martins (2015); Martins et al. (2016).

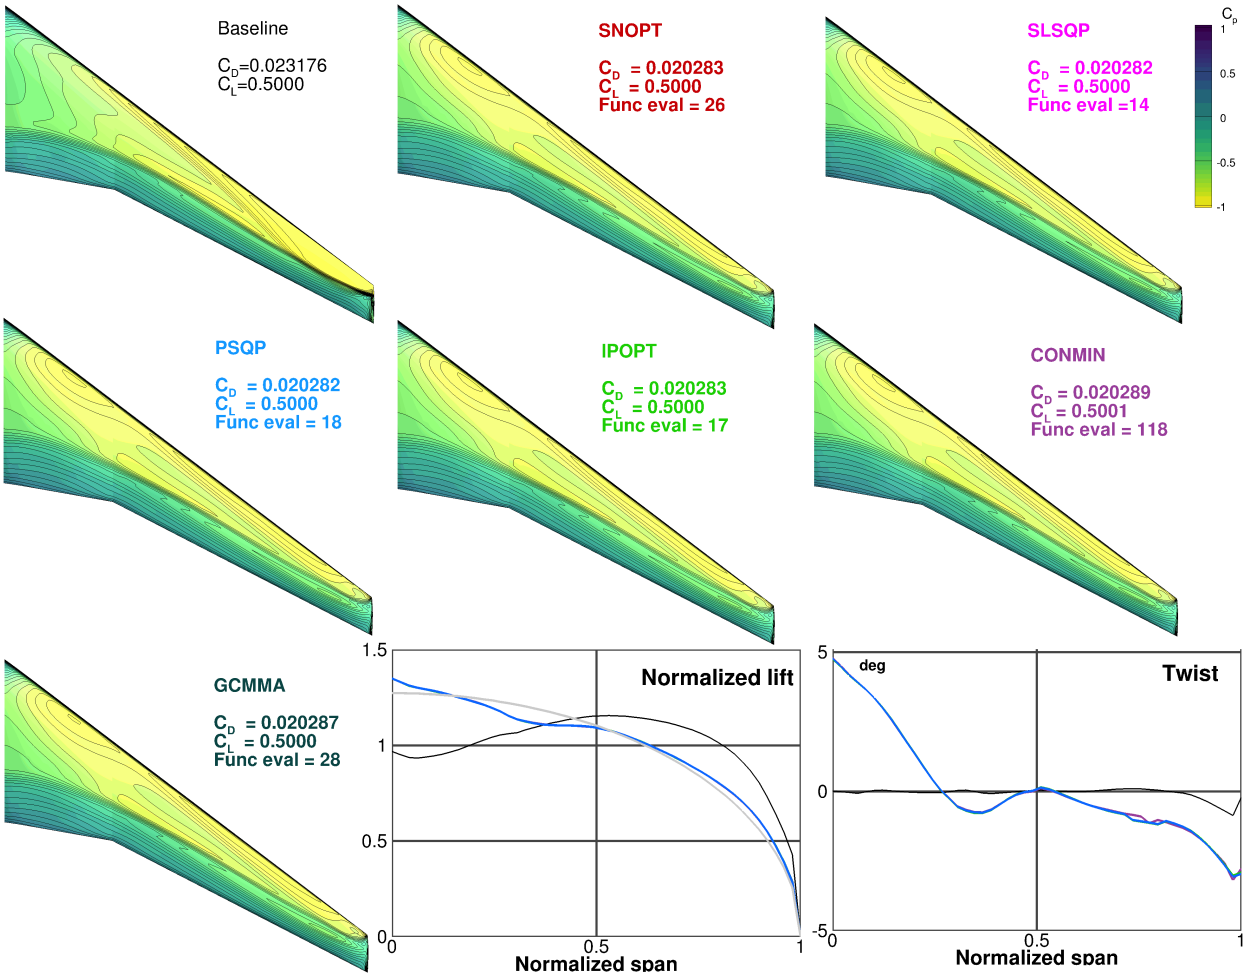

Figure 9: Aerodynamic twist optimization comparison on the L2 grid.

The baseline design has a strong shock, as evidenced by the closely spaced C_p contours, while the optimized designs have parallel, equally spaced pressure contours. All the optimized shapes are similar. The difference in the C_D values from different optimizers is within 1 drag count. The convergence history is shown in Figure 12; we see that SNOPT converges the fastest. A detailed comparison of the computational time is given in Table 4.2.

5 Influence of initial design

In the previous sections, we concluded that gradient-based optimization algorithms are more suitable for high-fidelity aerodynamic shape optimization problems because gradient-free algorithms require many more iterations to converge and thus incur an unacceptable computational cost for even a moderate number of design variables. However, gradient-based algorithms converge to a single minimum, and if the optimization problem is multimodal, these algorithms might not find the best possible design.

As mentioned in the introduction, Lyu et al. (2015) found a cluster of multiple local minima that were very close to each other. However, they tried only three different initial points, and

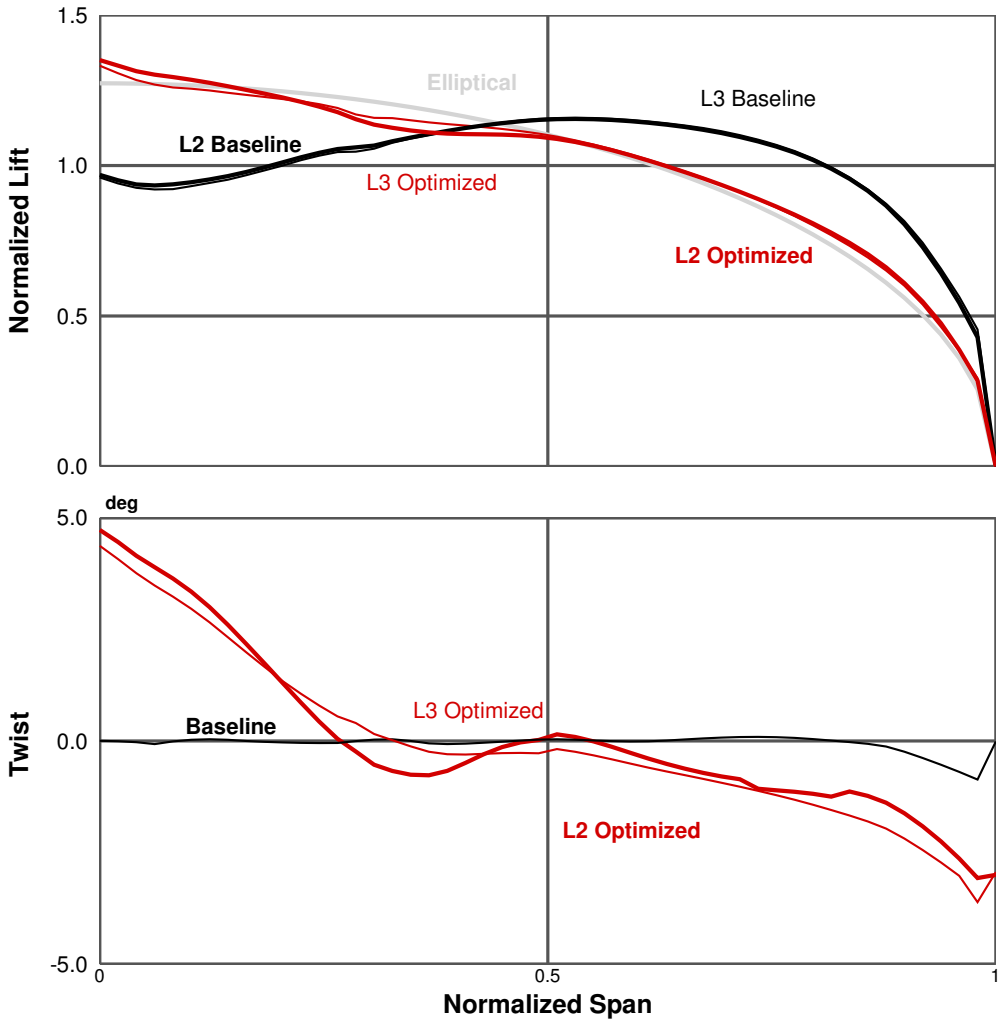

Figure 10: Lift and twist distribution comparison of the twist optimized designs.

so they could only speculate on the nature of the cluster. To explore this further, we solve the same ADODG CRM wing case starting from ten different geometries with combinations of random initial twist distributions and airfoil shapes. In addition, we examine the effect of tightening the optimality tolerance and refining the grid. We also solved the problem starting from a version of the CRM wing with NACA 0012 airfoils and zero twist to see if it converged to the same design.

5.1 CRM wing optimization benchmark

As previously mentioned, the results in this paper focus on the ADODG CRM wing optimization case Lyu et al. (2015), which does not include wing planform variables. In this section, we present the reference results for the single-point aerodynamic shape optimization of the CRM wing as defined by the ADODG. These results serve as a benchmark for the random initial cases that we present in the following sections. The CRM wing is optimized by minimizing the drag coefficient while enforcing both lift and moment constraints, and the CFD computations are performed on an L2 grid (see Table 3.2).

Figure 13 shows a detailed comparison of the CRM baseline wing (in red) and the optimized

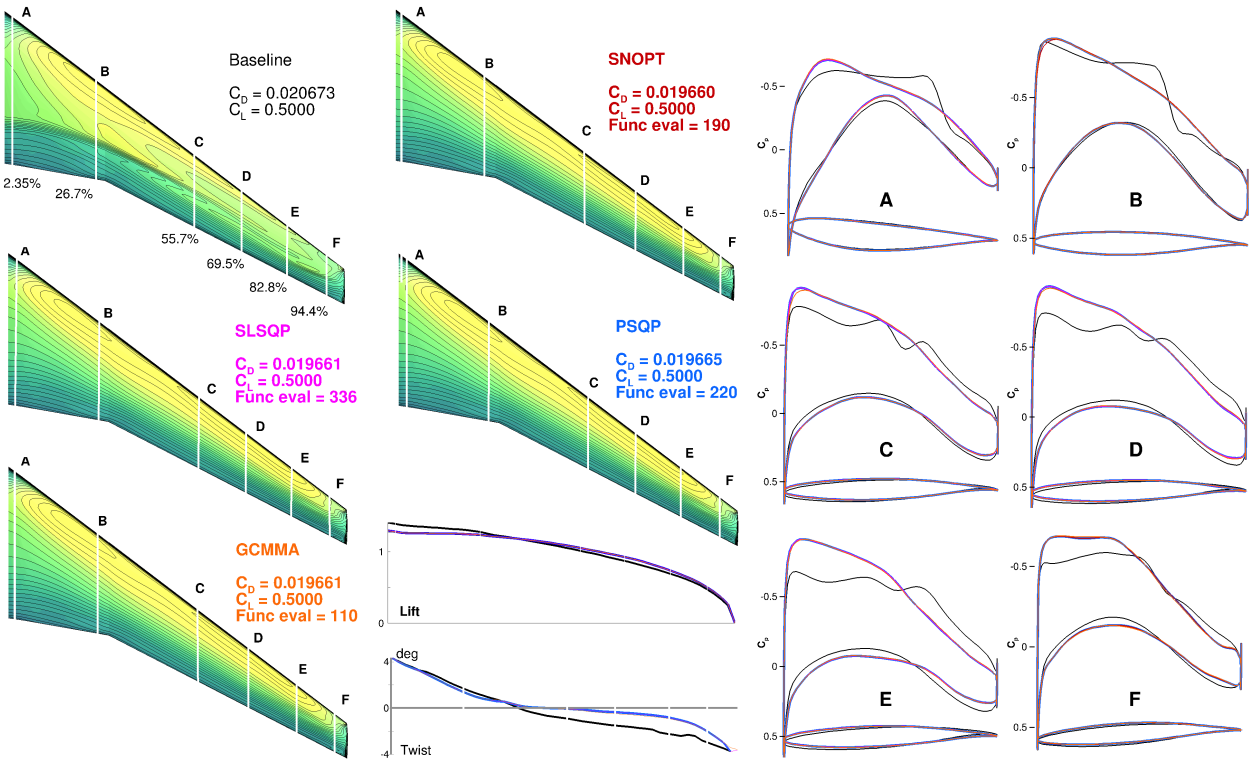

Figure 11: Comparison of optimizer performance in aerodynamic shape optimization.

Optimizer	Iterations	Proc hours
SNOPT	92	224.98
SLSQP	116	306.38
PSQP	221	562.60
GCCMA	298	772.60

Table 3: Computational cost comparison of the shape optimization on the L2 grid

wing (in blue). The optimization reduced the drag from 264.6 counts to 195.9 counts. At the optimum, the lift coefficient and the pitching moment coefficient are both at the target value. The lift distribution of the optimized wing is even closer to the elliptical distribution, and thus the induced drag of the optimized wing is closer to the theoretical minimum. The optimized wing exhibits more twist at the root and tip, and less twist in the mid-wing sections, while the baseline design has a near-linear twist distribution. The optimized thickness distribution is significantly different from the baseline. To reduce wave drag, the optimizer dramatically decreases the airfoil thickness on the outboard of the wing to the lower bound (25% of the original thickness), while increasing the thickness near the root to satisfy the volume constraint. These results match those of Lyu et al. (2015) for the single-point case with the L2 grid using 720 shape design variables. We use 720 shape design variables and 10 twist design variables, and the wing twist is no longer controlled solely by the shape design variables.

The optimizer satisfies the lift and moment constraints by changing both the angle of attack

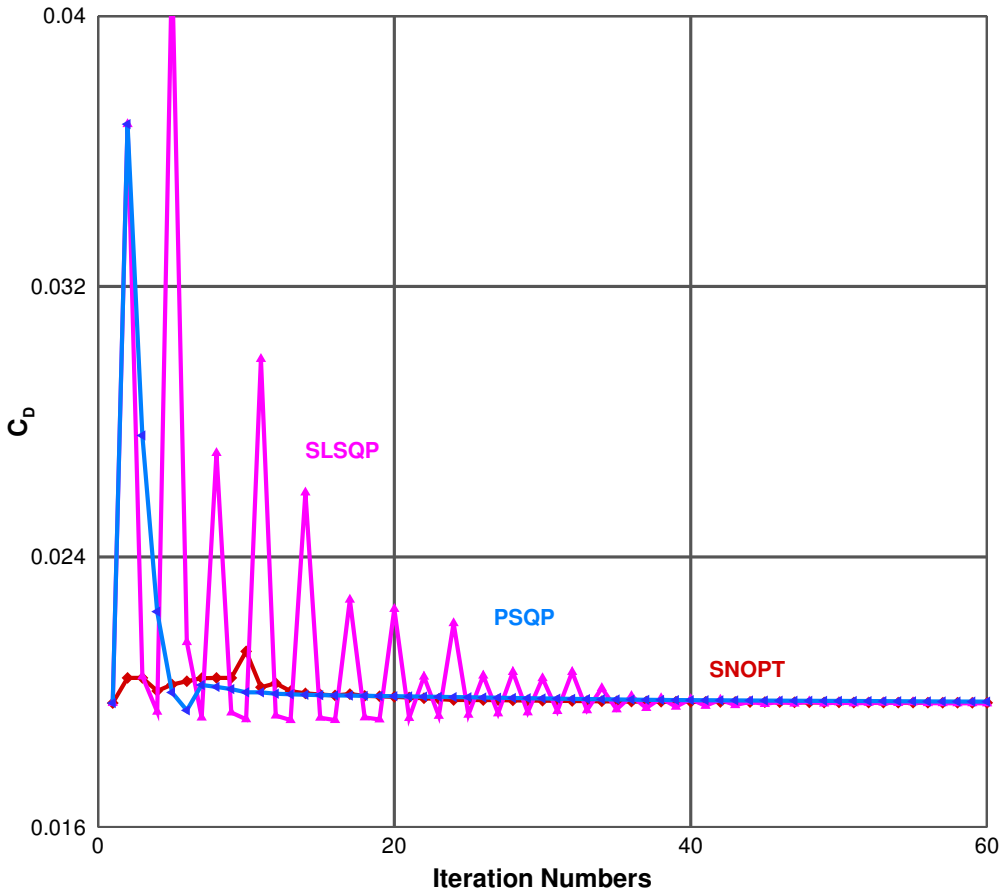

Figure 12: Convergence history of shape optimization for gradient-based methods on the L2 grid Lyu et al. (2014).

and the twist distribution. Without the twist design variables, the twist distribution is controlled only through the shape variables by movement of the FFD control points in the vertical direction, which can cause mesh failure if the optimal twist distribution differs too much from the initial one. An independent set of twist design variables that rotate the airfoils (thus preserving the chord lengths) makes it easier for the optimizer to reach the optimum twist, and this is crucial in cases such as those discussed in Section 5.8, where the initial wing twist is set to zero.

5.2 Random initial twist

As a first set of cases, we allow only the twist design variables to vary. The initial designs are generated by perturbing the twist design variables randomly. In Figure 15, we compare the results for the three initial random twist cases, where each optimized result is color-coded. Despite the poor initial geometries, the optimizer is able to smooth out the twist and achieve shock-free wings. All the constraints are met, and the lift distribution is close to elliptical.

Overall, the optimal geometries for the three initial points are very similar. The minimum drag for this case is much higher than that of the full benchmark case because only the twist design variables can be varied, and there is no way to shape the airfoils. Among the three random twist cases, the difference in drag is within 0.1 counts, and the difference in the C_p distribution is equally

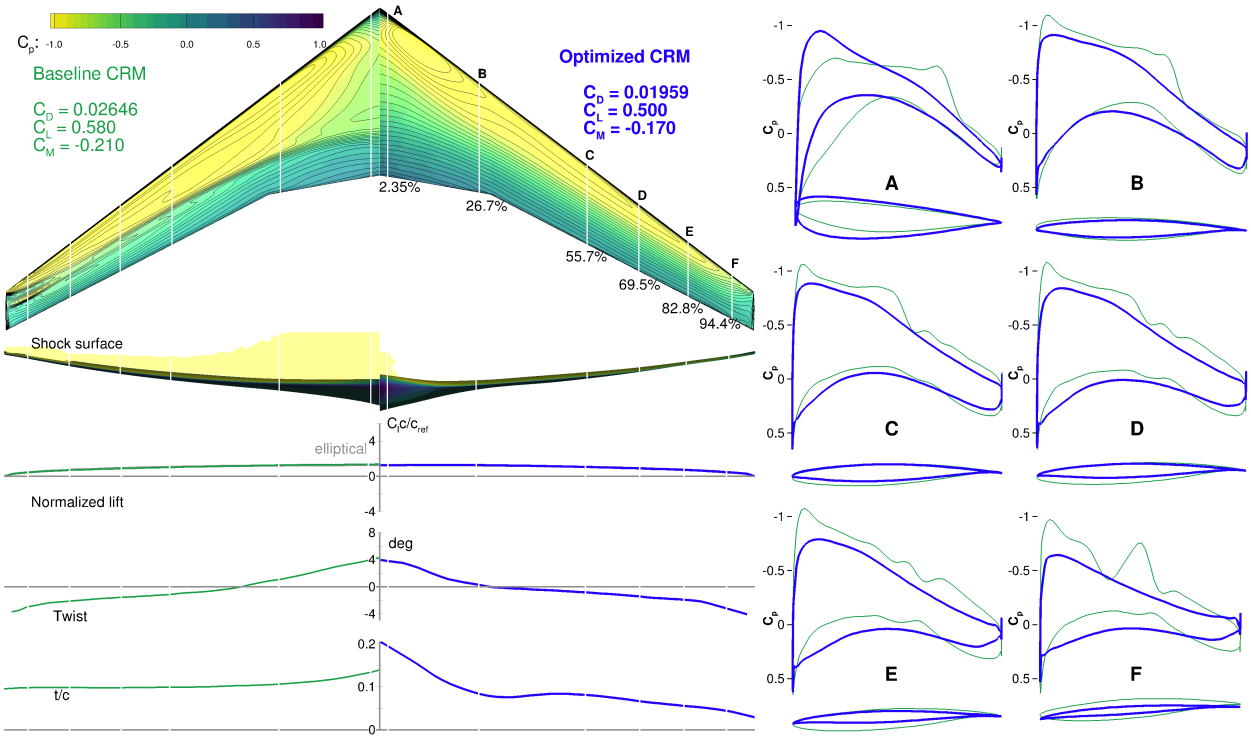

Figure 13: CRM wing optimization with shape and twist design variables.

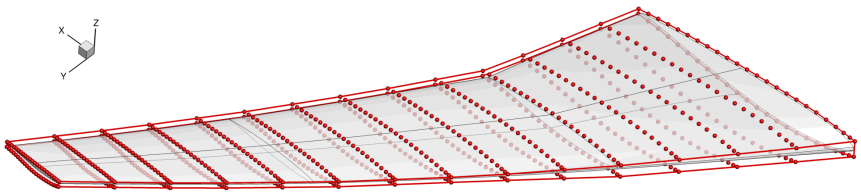

Figure 14: The number of FFD control points is increased to 720 for a more refined optimal shape.

negligible. The twist distributions for the three cases are also practically identical. These results support the hypothesis that there is only one global minimum in this problem.

5.3 Random initial shape

To further investigate the possibility of multimodality for this problem, we fix the twist design variables, leaving only the shape variables to be varied. Although the optimizer cannot change the twist distribution directly, it can do so indirectly through the airfoil shape variables by moving the FFD control points in the vertical direction, effectively shearing the airfoils.

The random perturbation on the airfoil thickness results in a even worse initial design than for the random twist cases. The jagged edges and wildly oscillating thickness distribution cause an equally chaotic C_p distribution (see the left-hand side of Figure 16). The RANS solution for this initial geometry is probably inaccurate because the flow separation is not modeled correctly. Despite the rather poor initial design and the likely inaccurate solution, the gradients point in

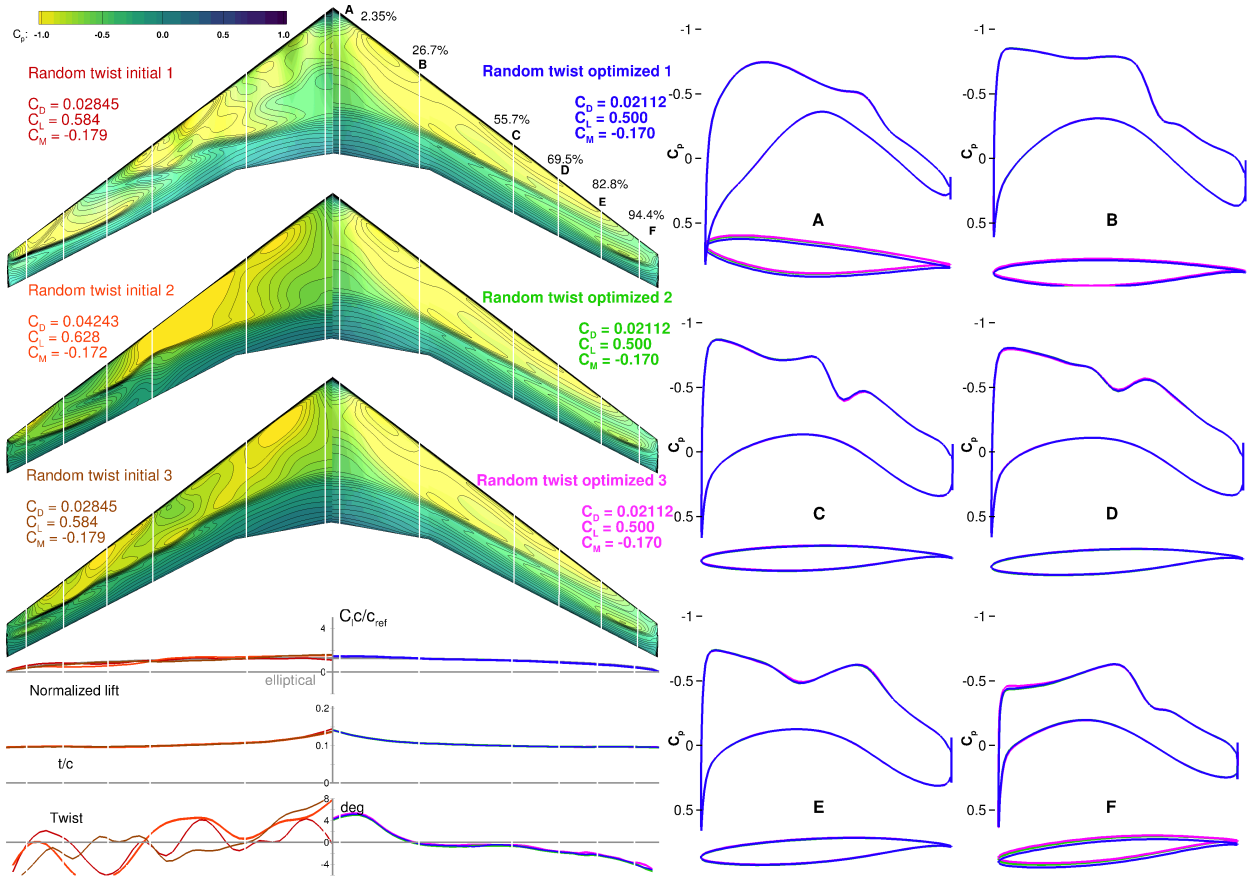

Figure 15: All three optimizations with random twist initial geometries converge to similar optimal designs.

the right direction, and the optimizer smooths out the wing and eventually achieves shock-free geometries.

Starting from wildly different initial designs, the three cases converge to similar optima, as shown in Figure 16. The lift and moment constraints are satisfied, and the drag count is only 0.4 counts higher than that of the benchmark case. We provide a more detailed comparison of a random initial case and the benchmark case in Section 5.4. Both the thickness distributions and the twist distributions in the optimized results are consistent among the optima. Small differences in the C_p distribution can be seen on the wing tips. The root and tip of the wing in case 3 are shifted compared to the other two cases. However, this has little effect on the aerodynamics, as shown in the C_p distributions. This indicates the possibility that the design space is multimodal.

5.4 Random initial twist and shape

We now use both twist and airfoil shape variables and apply random perturbations to both of these sets of design variables to produce the initial geometries. As shown in Figure 17, the optimizer converges to three extremely close geometries despite the great difference in the initial points. The differences in drag are less than 0.1 count, and the drag values are closer to that of the nominal

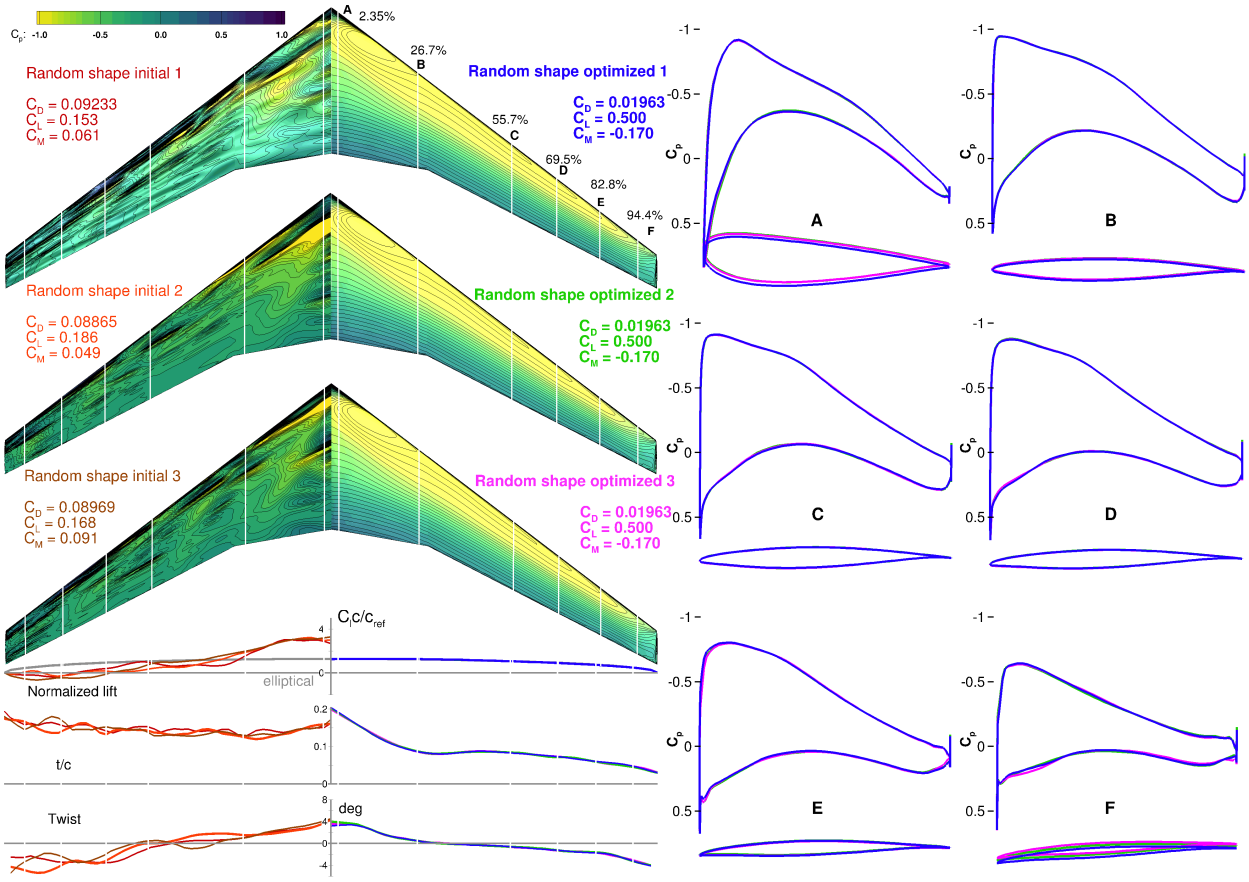

Figure 16: The optimizations starting from random shape variable perturbations converge to practically the same optimal shape.

case. All the optimized geometries have almost identical airfoil shape and twist distribution, with negligible differences in the C_p distribution. The shift in the root and tip of the wing in case 3 still exists, similarly to that shown in Figure 16.

A comparison of the second random initial case and the nominal CRM optimization case is shown in Figure 18. Overall, the two optimization results are very close to each other. The drag in the nominal case is 0.2 counts lower than that for the random initial case. In the root section, the nominal optimized wing is slightly thicker than that in the random initial case, resulting in a slightly better C_p distribution. The spanwise thickness and twist distributions are virtually identical.

To further investigate multimodality, we analyze the design space around the optimal points in more detail. Although it is impossible to fully visualize the 730-dimensional design space, we can visualize one-dimensional slices. We define these slices as the lines connecting the various optimal points, and we run the CFD for a series of geometries corresponding to points along these lines. For each point, we compute the corresponding optimization merit function value, as well as the objective function and constraint violation, as shown in Figure 19. The merit function value is a linear combination of the objective function and constraints Gill et al. (2007). A violation of the constraints is factored into the merit function calculation by the addition of a weighted penalty to the objective function value. As shown in Figure 19, between each pair of cases, the drag

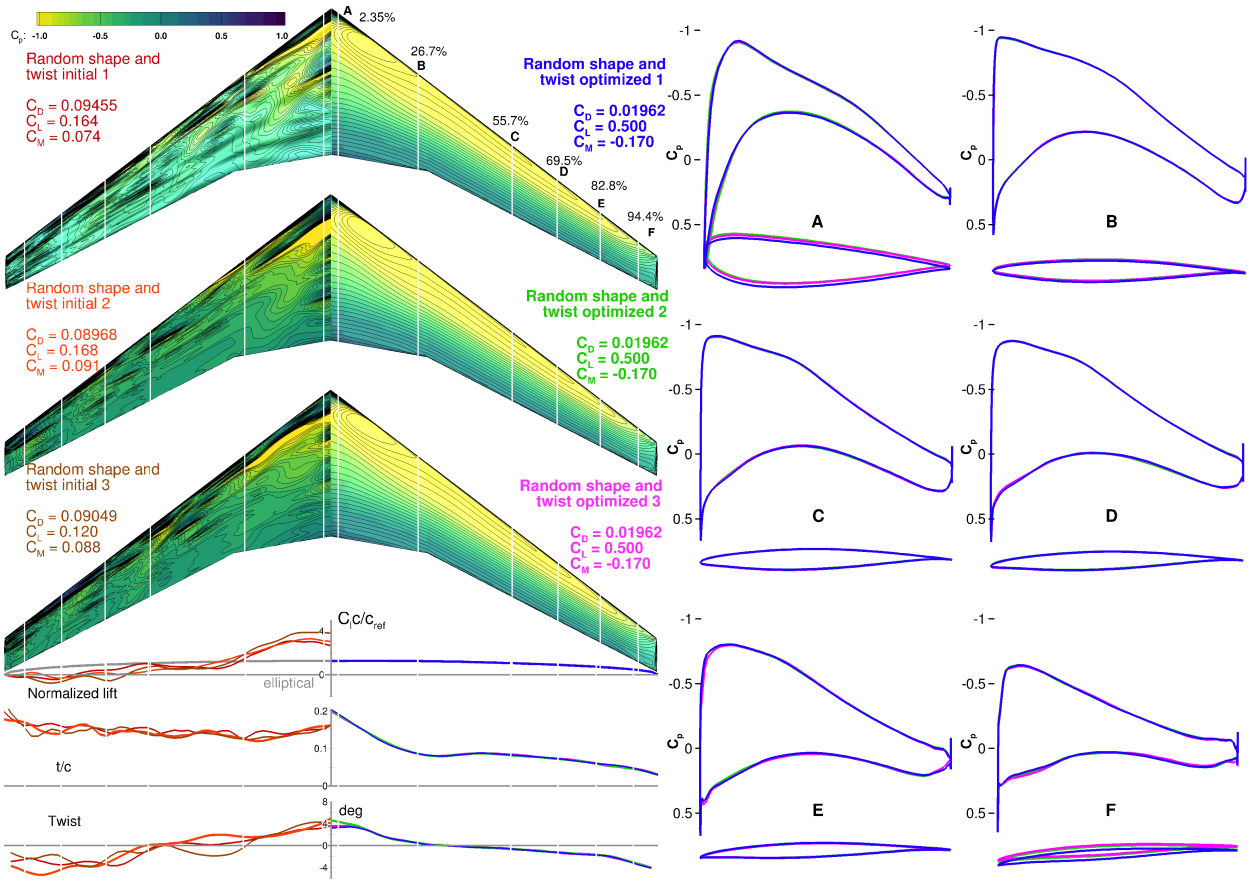

Figure 17: Comparison of the shape and twist optimization cases.

is monotonically increasing or decreasing. However, because of the weighted constraint violation penalty, the merit function has a local minimum at each optimized geometry.

We conclude from these results that each initial point converged to a different optimum. However, these optima are extremely close to each other, both in terms of the drag value and the value of the design variables. Our hypothesis is that this multimodality is not physical and is instead an artifact of the CFD numerical errors. In the next section, we verify this by solving the same problems with a smaller CFD convergence tolerance.

5.5 Effect of decreasing the convergence tolerance

To establish the nature of the multimodality observed in the previous section, we decrease the CFD convergence tolerance from 10^{-6} to 10^{-8} and the adjoint solver convergence tolerance from 10^{-10} to 10^{-12}, and we rerun the optimizations. Figure 20 shows that the minimum drag decreases by 0.1 counts. The twist distributions and C_p distributions are more alike than those shown in Figure 17. The optimizer is also able to further improve the airfoil shape at the root and tip sections, resulting in slightly different C_p distributions, while the vertical shift at the tip sections is reduced. The fact that these results are now closer together reinforces the conclusion that these multiple local minima are numerical rather than physical.

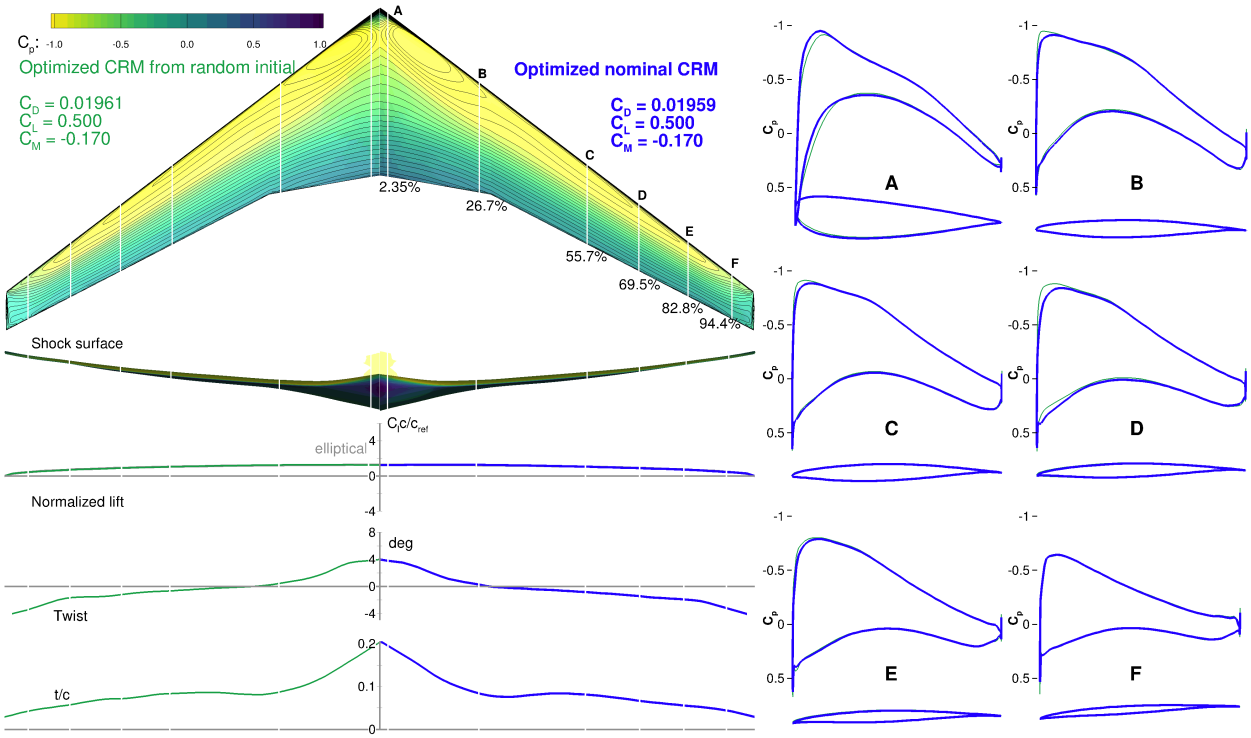

Figure 18: The nominal and random initial optimizations including shape and twist variables converge to practically the same optimum.

5.6 Effect of removing the angle-of-attack design variable

We now remove the angle of attack to establish if this design variable plays a role in the multimodality observed so far. We rerun the three random initial shape and twist optimization cases with a fixed angle-of-attack value of 2.8 deg. We choose this value because it is close to the value to which most of the optimization cases converge.

In the original ADODG case definition for aerodynamics shape optimization based on the CRM wing, the design variables include both twist and angle of attack. In the optimization cases described so far, we used both of these sets of design variables. The twist design variables rotate each section of the control points in the FFD volume about the corresponding reference point on the reference axis, while the angle-of-attack variable rotates all control points on the FFD volume about the same reference point. Following the ADODG case definition, the trailing edge of the wing and the root section are fixed, and only the leading edge of the wing is free to move. This limitation prevents the optimizer from converging to an unrealistically large angle of attack while compensating with negative twist on the wing.

However, due to the difference of definition between twist and angle of attack, even with the same local incidence values (angle of attack plus twist), the geometries can be different. We can see this in Figure 17, where the three sets of optimized results have very close objective function values but slightly different combinations of angle of attack and twist distribution, especially near the wing root section.

The results for cases with only shape and twist design variables are shown in Figure 21. The three random initial cases converged to geometries with the same drag, lift, and moment. The drag

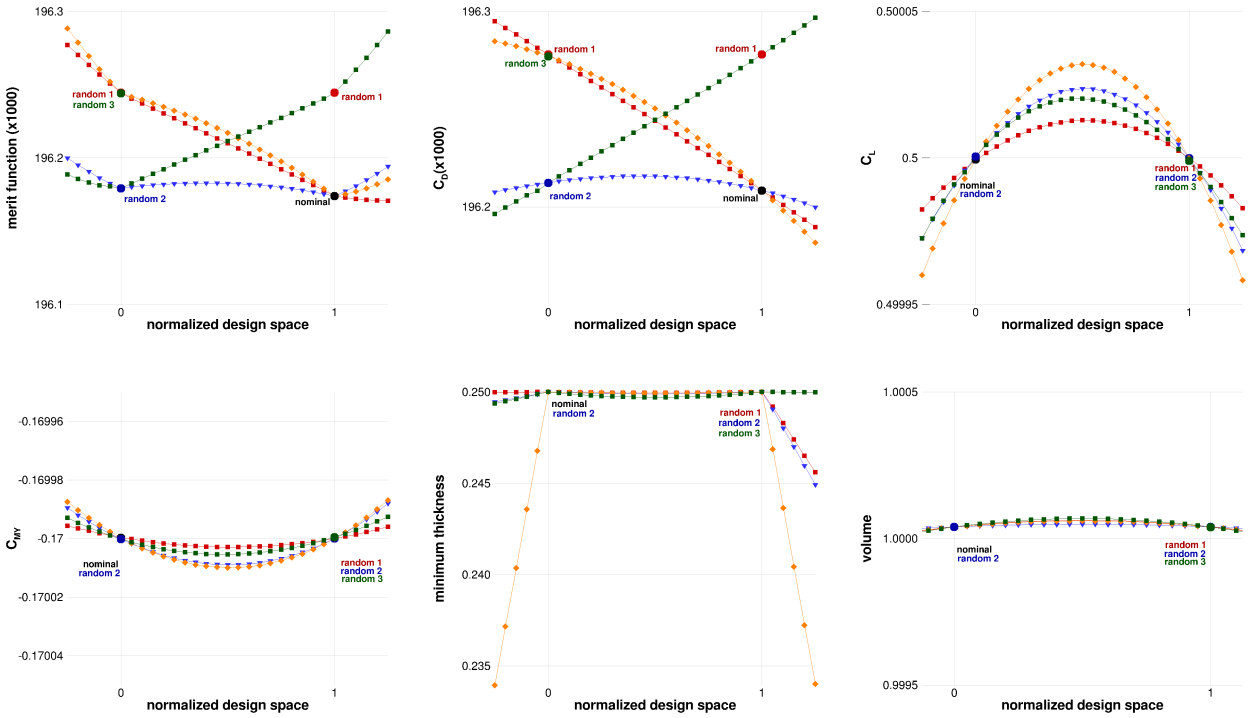

Figure 19: Merit function results and breakdown show the existence of numerical local minima; however, the differences in drag between the local minima are negligible.

values for these cases are very close to those seen in the cases where both angle of attack and twist are design variables. We conclude that removing the angle-of-attack variable does not decrease the drag. The twist distributions, however, are much closer compared to the previous cases, and the difference in twist at the root section is completely eliminated. This means that there exist multiple geometries with different combinations of angle of attack and twist that satisfy the optimization tolerance. Fixing the angle of attack saves the optimizer from navigating a design space with more flatness, which reduces the computational effort and the sources of local minima.

5.7 Effect of grid refinement

We now take the three optimized geometries and apply the same design variable values to a finer mesh (L1) with a FFD volume of the same size. The same optimization problem is then resumed on this finer mesh, starting from the geometries optimized using the coarser mesh (L2). A similar approach of restarting an optimization with finer meshes has reduced the computational cost in previous work Lyu et al. (2013). The optimization cases finish with identical optimized geometries, and the drag count is reduced by 11.55 counts. A similar trend of drag reduction when refining the mesh was seen by Lyu et al. (2015).

We perform the same one-dimensional slice analysis as in the L2 cases; see Figure 22. The L1 results show a similar trend of merit function change, but with a much smaller distance between each pair of optima. The largest change in the merit function value among the three cases is a drop by a factor of 5 (from 0.05% to 0.01%). The main reason is that the drags of the optimized geometries become much closer to each other on the finer mesh. This reinforces the idea that the cluster of local minima is due to the numerics, and that the radius of this cluster reduces as the

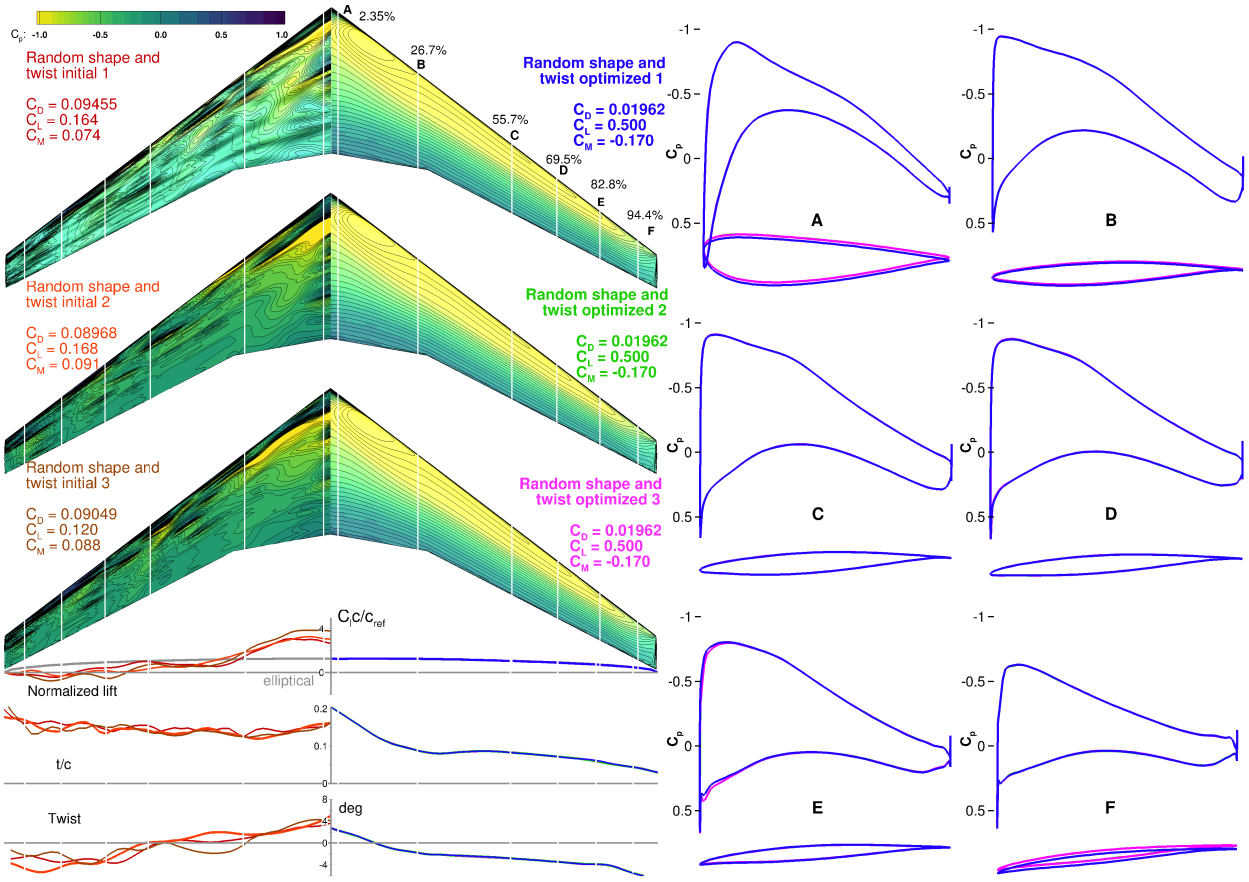

Figure 20: Comparison of the shape and twist optimization cases with better convergence.

convergence tolerance decreases and the grids are refined.

Earlier in the history of the ADODG CRM wing benchmark, other researchers found multiple solutions that were at the time assumed to be multiple local minima Dumont and Méheut (2016); Osusky (2013); LeDoux et al. (2015). However, initial studies by Lyu et al. (2015) and a subsequent result by Koo and Zingg (2018) contradict this assumption. The results presented herein provide a more detailed analysis that reinforces the conclusion that the design space for this problem is unimodal for practical purposes. We hypothesize that the other efforts that did not converge to the same result might not have converged the solver, the adjoint, or the optimization well enough to get results closer to ours. Accurate drag and lift computations, and accurate gradient computations that are consistent with the solver, are crucial because there is a subtle trade-off between thinning the outboard to reduce the drag, and thickening the inboard to meet the volume constraint, which contributes to a flat design space Lyu et al. (2015). This trade-off requires an effective and well-tuned optimizer, with appropriate design variable scaling, and an accurate and consistent adjoint solver.

5.8 CRM with zero initial twist and NACA 0012 airfoil

In the results we have seen so far, the initial designs were random perturbations of the baseline CRM geometry. However, these initial geometries are all based on the original CRM geometry,

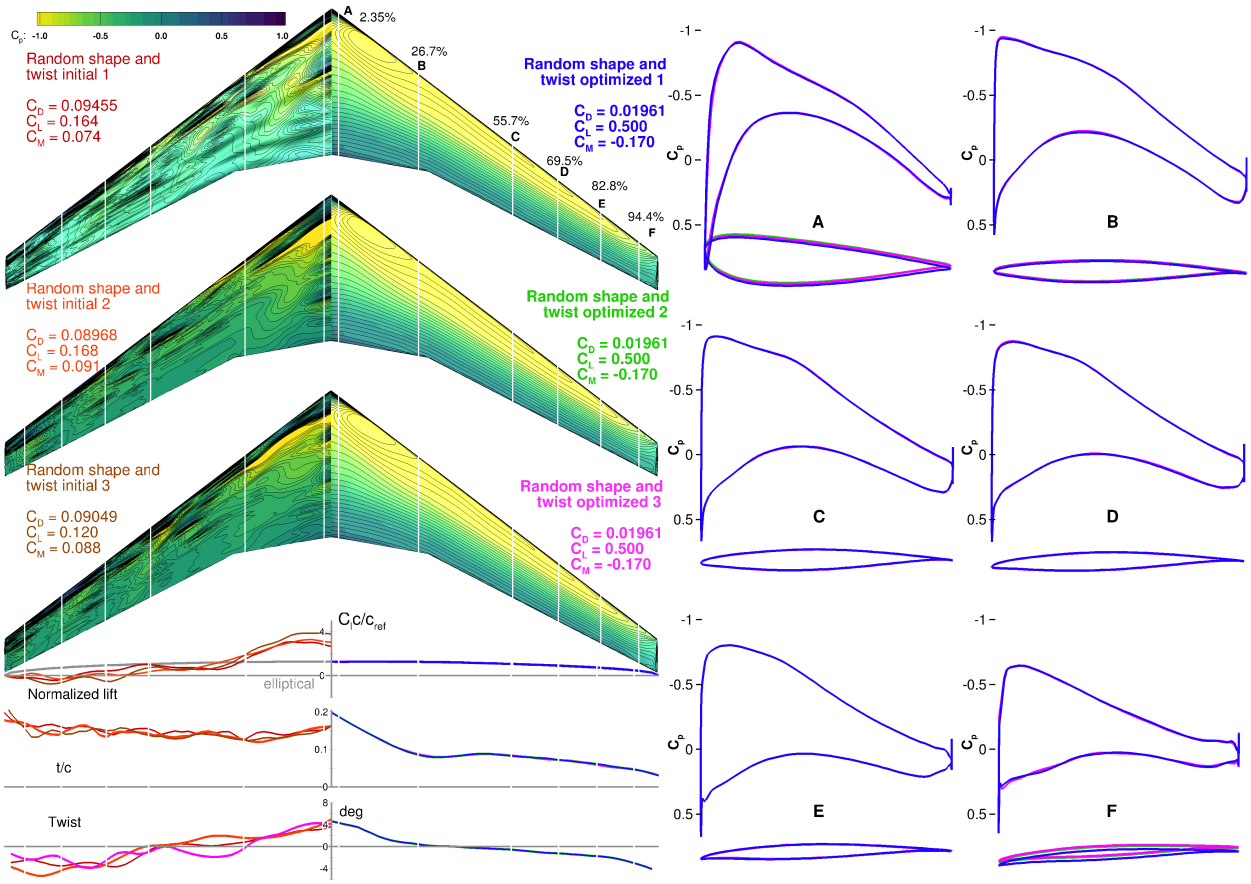

Figure 21: Optimizations without angle-of-attack design variable.

and therefore the results could be biased in some way by this common origin. To provide an initial design with no relation to the CRM, we construct a design that can be considered a "blank slate" that has no design rationale, but still looks like a wing. This design has the same planform as the CRM wing, but instead of the carefully designed airfoil stack and twist distribution of the CRM, it uses a NACA 0012 airfoil throughout and has constant twist and t/c distributions. We optimize starting from this geometry using 192 shape design variables, and then we compare the results with a CRM wing optimized using the same number of design variables.

Figure 23 shows that the initial NACA 0012 design exhibits a strong shock on the upper surface of the wing, even though the initial lift coefficient is just over half of what we ultimately require. Starting from this bad design, the optimizer is able to converge to a twist distribution similar to that of the original CRM, and it achieves a shock-free wing design for which all constraints are satisfied.

Figure 25 compares this result to the original optimization starting from the CRM. We can see only small differences in the C_p distribution and airfoil shape. The difference in drag is around 0.4 counts. Although the optimizations start from two completely different geometries, they end up close to each other, with only minor differences in the airfoil shape. There is a more significant difference in the dihedral, but this does not significantly affect the aerodynamic performance. As

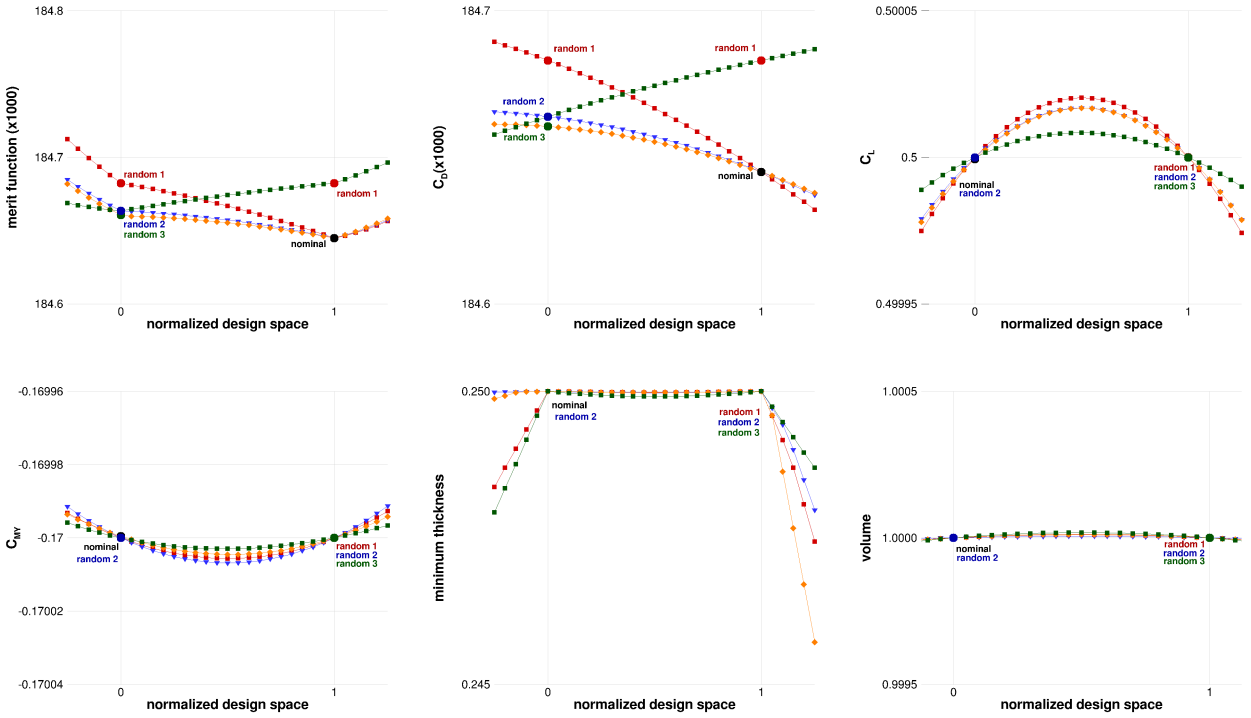

Figure 22: Merit function results and breakdown of optimization results using a finer grid.

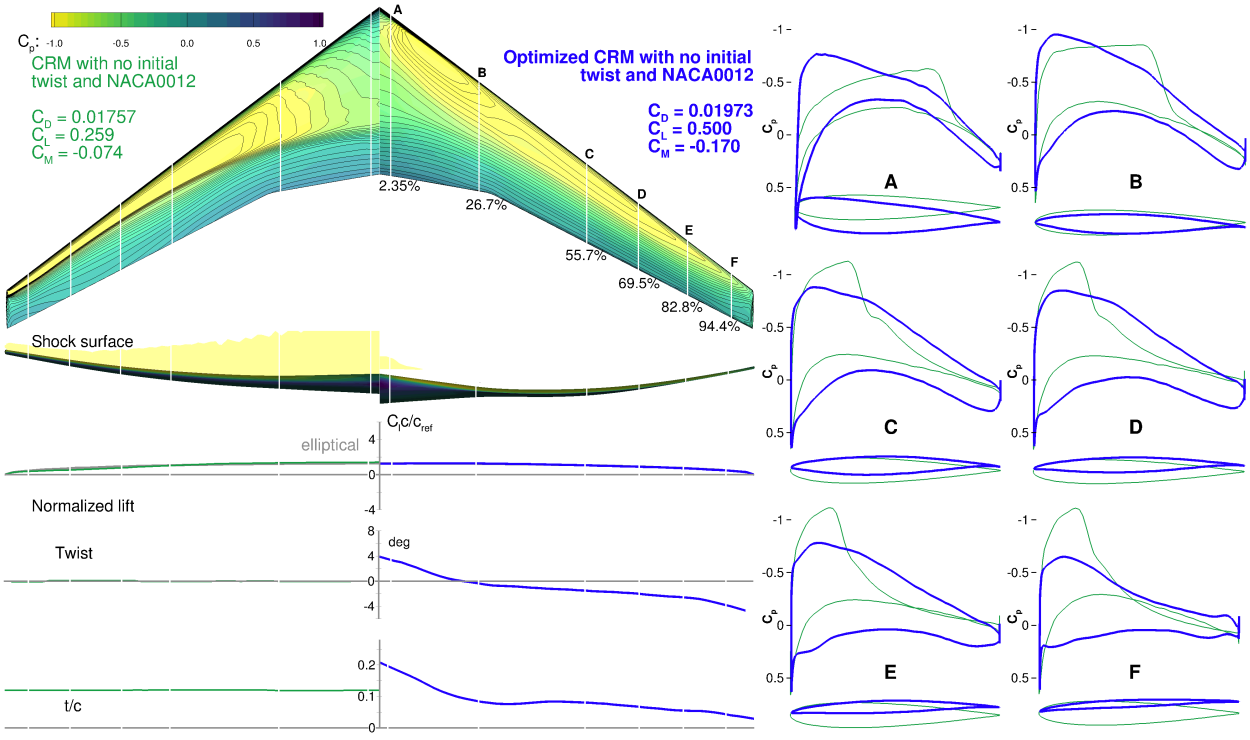

Figure 23: Optimized wing starting from CRM with NACA 0012 sections and no initial twist.

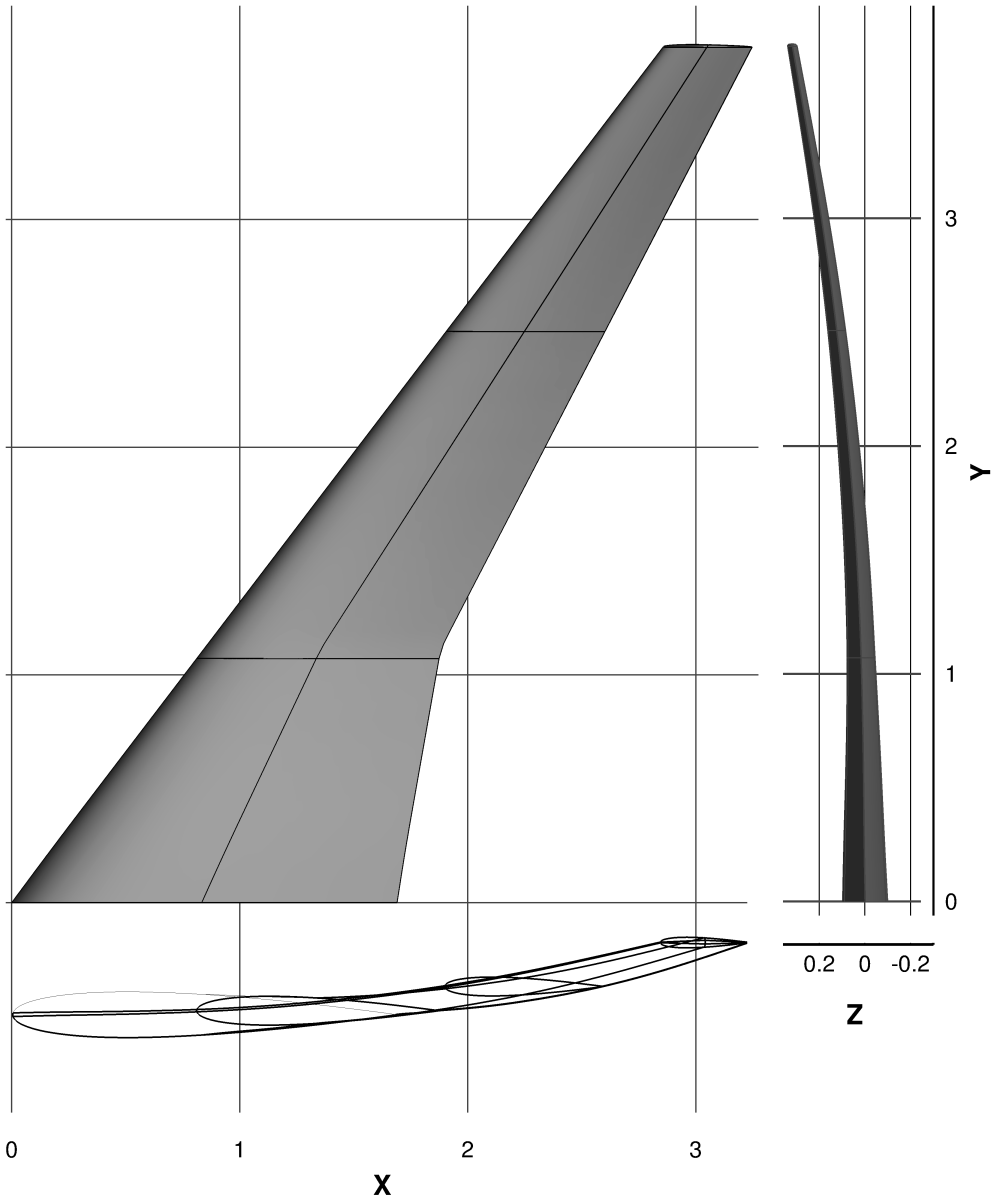

Figure 24: Baseline geometry of CRM with zero twist and NACA 0012 airfoil.

in the previous optimizations, this points us toward the conclusion that the aerodynamic shape optimization design space does not exhibit physically significant multimodality for a fixed wing planform.

6 Conclusions

In this paper, we have examined the effect of using different optimization algorithms and starting from different initial designs for aerodynamic shape optimizations of a fixed wing planform. To examine the effect of using different optimization algorithms, we benchmarked the performance of

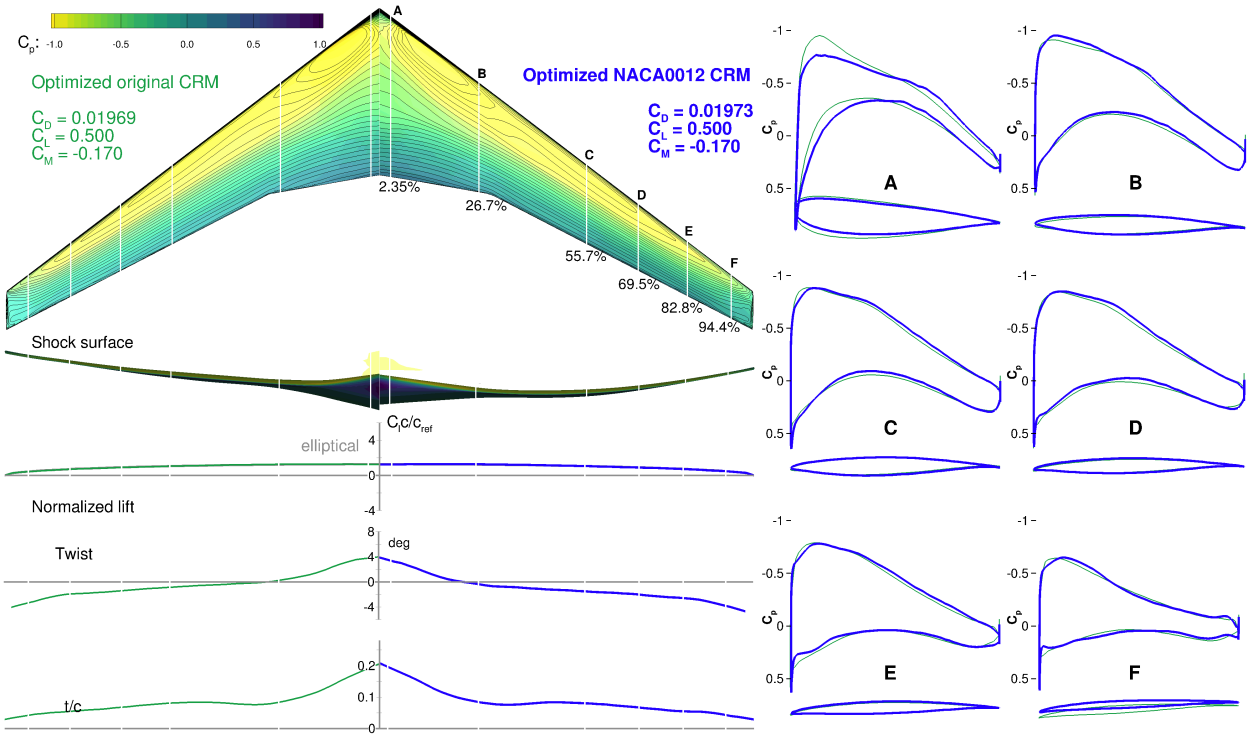

Figure 25: Comparison of the optimization starting from the CRM with the optimization starting from the NACA 0012 wing, showing that they find similar optima.

six gradient-based and three gradient-free optimizers for a single point wing aerodynamic shape optimization problem. To examine the effect of starting with different initial designs, we solved the same wing optimization problem using a gradient-based algorithm starting from a variety of designs.

For the optimization algorithm benchmarking we used only eight twist variables and the coarsest CFD grid because of the large number of evaluations required by the gradient-free optimization algorithms. All the optimizers except NSGA2 converged to the reference optimum, with identical twist distributions and only a 0.1 difference in the drag counts. The gradient-free methods were several times more expensive than the gradient-based methods, ranging from four times more expensive (NOMAD) to over 100 (NSGA2). Thus, gradient-based methods are a better choice because they require many fewer CFD evaluations and are equally robust.

We also benchmarked four of the gradient-based algorithms separately for a higher dimensional problem (192 design variables including twist and airfoil shape) using a finer grid. All the optimizers converged to the same optimum, reducing the drag by 4.84%, from 206.7 to 196.6 counts. SNOPT converged the fastest, using just under 225 processor-hours.

To examine the effect of starting from different initial designs, and the associated issue of multimodality, we solved the wing aerodynamic shape optimization problem using SNOPT for the medium grid (450 thousand cells) and up to 720 shape variables and 10 twist variables. We solved the 14 cases listed in Table 5.8. The resulting shapes were practically identical, with drag values within 0.1 counts.

Although these optima were the same for practical purposes, they did represent multiple local minima from a numerical point of view. To analyze this phenomenon in more detail, we studied

Initial shape	Drag (counts)
CRM	195.9
Random twist 1	211.2
Random twist 2	211.2
Random twist 3	211.2
Random shape 1	196.3
Random shape 2	196.3
Random shape 3	196.3
Random twist and shape 1	196.2
Random twist and shape 2	196.2
Random twist and shape 3	196.2
Random twist and shape, no AoA 1	196.1
Random twist and shape, no AoA 2	196.1
Random twist and shape, no AoA 3	196.1
Untwisted NACA 0012	197.3

Table 4: Summary of optimization results

the effect of decreasing the convergence tolerances and refining the CFD grids. We observed that the differences between the local minima decreased, further supporting the idea that they are due to discretization and convergence limitations.

Finally, we solved the same optimization problem starting from a geometry that has the CRM wing planform but with zero twist, constant t/c, and a NACA 0012 airfoil. We optimized the wing with respect to 192 shape variables and 10 twist variables using a 450K-cell grid. The result of this optimization was close to that of the nominal CRM optimization, with minor differences in airfoil shape.

Overall, these results support the conclusion that the design space for wing design optimization with a fixed problem is largely convex, with a small flat region that is multimodal. However, this region is so small, and the differences in drag so minor, that the design space can be considered unimodal for all practical purposes.

7 Acknowledgments

The authors would like to thank Gaetan K. W. Kenway for the mesh generator and for his support of this project, and Nicolas Bons for assisting with the NOMAD simulation results, as well as our colleagues at MDO Lab for the numerous ideas. The computations were performed on the Stampede HPC of the Extreme Science and Engineering Discovery Environment (XSEDE) and on the Flux HPC cluster at the University of Michigan Advanced Computing Center.

8 Bibliography

References

Z. Lyu, J. R. R. A. Martins, Strategies for solving high-fidelity aerodynamic shape optimization problems, in: 15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Atlanta, GA, 2014. doi:doi:10.2514/6.2014-2594, AIAA 2014-2594.

A. Jameson, Aerodynamic shape optimization using the adjoint method, Lecture Series, Von Karman Institute for Fluid Dynamics, Rode Saint Genèse, Belgium, 2003.

R. M. Hicks, P. A. Henne, Wing design by numerical optimization, Journal of Aircraft 15 (1978) 407–412.

J. E. V. Peter, R. P. Dwight, Numerical sensitivity analysis for aerodynamic optimization: A survey of approaches, Computers and Fluids 39 (2010) 373–391.

J. R. R. A. Martins, J. T. Hwang, Review and unification of methods for computing derivatives of multidisciplinary computational models, AIAA Journal 51 (2013) 2582–2599.

O. Pironneau, On optimum profiles in Stokes flow, Journal of Fluid Mechanics 59 (1973) 117–128.

O. Pironneau, On optimum design in fluid mechanics, Journal of Fluid Mechanics 64 (1974) 97–110.

A. Jameson, Aerodynamic design via control theory, Journal of Scientific Computing 3 (1988) 233–260.

Z. Lyu, G. K. Kenway, C. Paige, J. R. R. A. Martins, Automatic differentiation adjoint of the Reynolds-averaged Navier–Stokes equations with a turbulence model, in: 21st AIAA Computational Fluid Dynamics Conference, San Diego, CA, 2013. doi:doi:10.2514/6.2013-2581.

Z. Lyu, G. K. W. Kenway, J. R. R. A. Martins, Aerodynamic shape optimization investigations of the Common Research Model wing benchmark, AIAA Journal 53 (2015) 968–985.

A. Jameson, L. Martinelli, N. A. Pierce, Optimum aerodynamic design using the Navier–Stokes equations, Theoretical and Computational Fluid Dynamics 10 (1998) 213–237.

W. K. Anderson, V. Venkatakrishnan, Aerodynamic design optimization on unstructured grids with a continuous adjoint formulation, Computers and Fluids 28 (1999) 443–480.

E. J. Nielsen, W. K. Anderson, Aerodynamic design optimization on unstructured meshes using the Navier–Stokes equations, AIAA Journal 37 (1999) 1411–1419.

R. P. Dwight, J. Brezillon, Efficient and robust algorithms for solution of the adjoint compressible Navier–Stokes equations with applications, International Journal for Numerical Methods in Fluids 60 (2009) 365–389.

J. Brezillon, R. P. Dwight, Applications of a discrete viscous adjoint method for aerodynamic shape optimisation of 3D configurations, CEAS Aeronautical Journal 3 (2012) 25–34.

G. K. W. Kenway, J. R. R. A. Martins, Buffet onset constraint formulation for aerodynamic shape optimization, AIAA Journal 55 (2017) 1930–1947.

R. P. Liem, J. R. R. A. Martins, G. K. Kenway, Expected drag minimization for aerodynamic design optimization based on aircraft operational data, Aerospace Science and Technology 63 (2017) 344–362.

S. Chen, Z. Lyu, G. K. W. Kenway, J. R. R. A. Martins, Aerodynamic shape optimization of the Common Research Model wing-body-tail configuration, Journal of Aircraft 53 (2016) 276–293.

G. K. W. Kenway, J. R. R. A. Martins, Multipoint aerodynamic shape optimization investigations of the Common Research Model wing, AIAA Journal 54 (2016) 113–128.

G. K. W. Kenway, J. R. R. A. Martins, Multipoint aerodynamic shape optimization investigations of the Common Research Model wing, in: Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech), Kissimmee, FL, 2015. doi:doi:10.2514/6.2015-0264.

K. Telidetzki, L. Osusky, D. W. Zingg, Application of jetstream to a suite of aerodynamic shape optimization problems, in: 52nd Aerospace Sciences Meeting, 2014. doi:doi:10.2514/6.2014-0571.

A. Elham, Adjoint quasi-three-dimensional aerodynamic solver for multi-fidelity wing aerodynamic shape optimization, Aerospace Science and Technology 41 (2015) 241–249.

M. Drela, Frontiers of Computational Fluid Dynamics, World Scientific, 1998, pp. 363–381. doi:doi:10.1142/9789812815774_0019.

M. Drela, Design and optimization method for multi-element airfoils aerospace design conference, in: Proceedings of the AIAA/AHS/ASEE Aerospace Design Conference, AIAA 1993-0969, Irvine, CA, 1993. doi:doi:10.2514/6.1993-969.

D. I. Jones, J. W. Finch, Comparison of optimization algorithms, International Journal of Control 40 (1984) 747–761.

N. Marco, S. Lanteri, J.-A. Désidéri, B. Mantel, J. Périaux, Parallel genetic algorithms applied to optimum shape design in aeronautics, in: C. Lengauer, M. Griebl, S. Gorlatch (Eds.), Euro-Par'97 Parallel Processing, Springer, Berlin, Heidelberg, 1997, pp. 856–863. doi:doi:10.1007/BFb0002826.

Y. He, R. K. Agarwal, Shape optimization of NREL S809 airfoil for wind turbine blades using a multi-objective genetic algorithm, in: Proceedings of the 32nd AIAA Applied Aerodynamics Conference, 2014. doi:doi:10.2514/6.2014-2845.

D. W. Zingg, M. Nemec, T. H. Pulliam, A comparative evaluation of genetic and gradient-based algorithms applied to aerodynamic optimization, European Journal of Computational Mechanics 17 (2008) 103–126.

S. Obayashi, T. Tsukahara, Comparison of optimization algorithms for aerodynamic shape design, AIAA Journal 35 (1997) 1413–1415.

T. L. Holst, T. H. Pulliam, Aerodynamic shape optimization using a real-number-encoded genetic algorithm, in: Proceedings of the 19th Applied Aerodynamics Conference, AIAA 2001-2473, Anaheim, CA, 2001.

D. Sasaki, M. Morikawa, S. Obayashi, K. Nakahashi, Aerodynamic shape optimization of supersonic wings by adaptive range multiobjective genetic algorithms, in: Evolutionary Multi-Criterion Optimization, 2001, pp. 639–652. doi:doi:10.1007/3-540-44719-9_45.

O. Chernukhin, D. W. Zingg, Multimodality and global optimization in aerodynamic design, AIAA Journal 51 (2013) 1342–1354.

N. P. Bons, X. He, C. A. Mader, J. R. R. A. Martins, Multimodality in aerodynamic wing design optimization, in: 18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, 2017.

A. Dumont, M. Méheut, Gradient-based optimization of CRM wing-alone and wing-body-tail configurations by RANS adjoint technique, in: 54th AIAA Aerospace Sciences Meeting, American Institute of Aeronautics and Astronautics (AIAA), 2016. doi:doi:10.2514/6.2016-1293.

C. Lee, D. Koo, K. Telidetzki, H. Buckley, H. Gagnon, D. W. Zingg, Aerodynamic shape optimization of benchmark problems using jetstream, in: Proceedings of the 53rd AIAA Aerospace Sciences Meeting, American Institute of Aeronautics and Astronautics (AIAA), 2015. doi:doi:10.2514/6.2015-0262.

D. Shi-Dong, C.-H. Chen, S. Nadarajah, Adjoint-based aerodynamic optimization of benchmark CRM wing, in: 35th AIAA Applied Aerodynamics Conference, American Institute of Aeronautics and Astronautics, 2017. doi:doi:10.2514/6.2017-3755.

D. Koo, D. W. Zingg, Investigation into aerodynamic shape optimization of planar and nonplanar wings, AIAA Journal 56 (2018) 250–263.

G. K. W. Kenway, G. J. Kennedy, J. R. R. A. Martins, Scalable parallel approach for high-fidelity steady-state aeroelastic analysis and derivative computations, AIAA Journal 52 (2014) 935–951.

G. K. W. Kenway, J. R. R. A. Martins, Multipoint high-fidelity aerostructural optimization of a transport aircraft configuration, Journal of Aircraft 51 (2014) 144–160.

G. K. Kenway, G. J. Kennedy, J. R. R. A. Martins, A CAD-free approach to high-fidelity aerostructural optimization, in: Proceedings of the 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference, AIAA 2010-9231, Fort Worth, TX, 2010. doi:doi:10.2514/6.2010-9231.

C. Lee, D. Koo, D. W. Zingg, Comparison of b-spline surface and free-form deformation geometry control for aerodynamic optimization, AIAA Journal 55 (2017) 228–240.

E. van der Weide, G. Kalitzin, J. Schluter, J. J. Alonso, Unsteady turbomachinery computations using massively parallel platforms, in: Proceedings of the 44th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, 2006. AIAA 2006-0421.

A. Jameson, W. Schmidt, E. Turkel, Numerical Solution of the Euler Equations by Finite Volume Methods Using Runge–Kutta Time Stepping Schemes, Technical Report AIAA 1981-1259, 1981.

C. A. Mader, G. Kenway, J. R. R. A. Martins, Towards high-fidelity aerostructural optimization using a coupled ADjoint approach, in: Proceedings of the 12th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Victoria, BC, 2008. AIAA 2008-5968.

J. R. R. A. Martins, G. K. W. Kenway, T. R. Brooks, Multidisciplinary Design Optimization of Aircraft Configurations—Part 2: High-fidelity aerostructural optimization, Lecture Series, Von Karman Institute for Fluid Dynamics, Rode Saint Genèse, Belgium, 2016. ISSN0377-8312.

J. R. R. A. Martins, P. Sturdza, J. J. Alonso, The complex-step derivative approximation, ACM Transactions on Mathematical Software 29 (2003) 245–262. September.

Y. Saad, M. H. Schultz, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, SIAM Journal on Scientific and Statistical Computing 7 (1986) 856–869.

S. Balay, W. D. Gropp, L. C. McInnes, B. F. Smith, Efficient management of parallelism in object oriented numerical software libraries, in: E. Arge, A. M. Bruaset, H. P. Langtangen (Eds.), Modern Software Tools for Scientific Computing, Birkhäuser Press, 1997, pp. 163–202.

S. Balay, J. Brown, K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, B. F. Smith, H. Zhang, PETSc Users Manual, Technical Report ANL-95/11 - Revision 3.4, Argonne National Laboratory, 2013a.

S. Balay, J. Brown, K. Buschelman, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, B. F. Smith, H. Zhang, PETSc Web page, 2013b. Http://www.mcs.anl.gov/petsc.

C. A. Mader, J. R. R. A. Martins, Stability-constrained aerodynamic shape optimization of flying wings, Journal of Aircraft 50 (2013) 1431–1449.

Z. Lyu, J. R. R. A. Martins, Aerodynamic design optimization studies of a blended-wing-body aircraft, Journal of Aircraft 51 (2014) 1604–1617.

R. P. Liem, G. K. W. Kenway, J. R. R. A. Martins, Multimission aircraft fuel burn minimization via multipoint aerostructural optimization, AIAA Journal 53 (2015) 104–122.

R. E. Perez, P. W. Jansen, J. R. R. A. Martins, pyOpt: A Python-based object-oriented framework for nonlinear constrained optimization, Structural and Multidisciplinary Optimization 45 (2012) 101–118.

P. E. Gill, W. Murray, M. A. Saunders, SNOPT: An SQP algorithm for large-scale constrained optimization, SIAM Journal of Optimization 12 (2002) 979–1006.

D. Kraft, A software package for sequential quadratic programming, Technical Report, Tech. Rep. DFVLR-FB 88-28, DLR German Aerospace Center, 1988. URL: `http://www.opengrey.eu/item/display/10068/147127`.

C. L. Lawson, R. J. Hanson, Solving Least Squares Problems, volume 15 of *Classics in Applied Mathematics*, SIAM, 1987. doi:doi:10.1137/1.9781611971217.

D. Kraft, Algorithm 733; TOMP—Fortran modules for optimal control calculations, ACM Transactions on Mathematical Software 20 (1994) 262–281.

A. Waechter, L. T. Biegler, On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming, Mathematical Programming 106 (2006) 25–57.

R. Fletcher, S. Leyffer, Nonlinear programming without a penalty function, Mathematical Programming 269 (2002) 239–269.

G. N. Vanderplaats, CONMIN user's manual, 1978.

K. Svanberg, The method of moving asymptotes—a new method for structural optimization, Numerical Methods in Engineering 24 (1987) 359–373.

P. Jansen, R. Perez, Constrained structural design optimization via a parallel augmented Lagrangian particle swarm optimization approach, Computers & Structures 89 (2011) 1352–1366.

R. Perez, K. Behdinan, Particle swarm approach for structural design optimization, International Journal of Computer and Structures 85 (2007) 1579–1588.

S. Haghighat, J. R. R. A. Martins, H. H. T. Liu, Aeroservoelastic design optimization of a flexible wing, Journal of Aircraft 49 (2012) 432–443.

P. Jansen, R. E. Perez, J. R. R. A. Martins, Aerostructural optimization of nonplanar lifting surfaces, Journal of Aircraft 47 (2010) 1491–1503.

K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II, IEEE Transactions on Evolutionary Computation 6 (2002) 182–197.

C. Audet, A. L. Custódio, J. E. Dennis, Mesh adaptive direct search algorithms for constrained optimization, SIAM Journal on Optimization 17 (2006) 188–217.

S. Le Digabel, C. Tribes, C. Audet, NOMAD User Guide Version 3.5.1, Group for Research in Decision Analysis, Montreal, Quebec, Canada, 2012.

S. Le Digabel, Algorithm 909: NOMAD: Nonlinear optimization with the MADS algorithm, ACM Transactions on Mathematical Software 37 (2011) 44:1–44:15.

J. C. Vassberg, M. A. DeHaan, S. M. Rivers, R. A. Wahls, Development of a common research model for applied CFD validation studies, 2008. AIAA 2008-6919.

Z. Lyu, Z. Xu, J. R. R. A. Martins, Benchmarking optimization algorithms for wing aerodynamic design optimization, in: Proceedings of the 8th International Conference on Computational Fluid Dynamics, Chengdu, Sichuan, China, 2014. ICCFD8-2014-0203.

G. J. Kennedy, G. K. W. Kenway, J. R. R. A. Martins, High aspect ratio wing design: Optimal aerostructural tradeoffs for the next generation of materials, in: Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech), National Harbor, MD, 2014. doi:doi:10.2514/6.2014-0596.

G. W. K. Kenway, J. R. R. A. Martins, High-fidelity aerostructural optimization considering buffet onset, in: Proceedings of the 16th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Dallas, TX, 2015. AIAA 2015-2790.

P. E. Gill, W. Murray, M. A. Saunders, User's Guide for SNOPT Version 7, A Fortran Package for Large-Scale Nonlinear Programming, Systems Optimization Laboratory, Stanford University, Stanford, CA 94305, 2007.

L. M. Osusky, A Novel Numerical Tool for Aerodynamic Shape Optimization in Turbulent Flow, Ph.D. thesis, University of Toronto, 2013.

S. T. LeDoux, J. C. Vassberg, D. P. Young, S. Fugal, D. Kamenetskiy, W. P. Huffman, R. G. Melvin, M. F. Smith, Study based on the AIAA aerodynamic design optimization discussion group test cases, AIAA Journal 53 (2015) 1910–1935.