

Multimodality in Aerodynamic Wing Design Optimization

Nicolas P. Bons¹, Xiaolong He², Charles A. Mader³, and Joaquim R. R. A. Martins⁴

Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI, 48109

Abstract The application of gradient-based optimization to wing design could potentially reveal revolutionary new wing concepts. Giving the optimizer the freedom to discover novel wing designs may increase the likelihood of multimodality in the design space. To address this issue, we investigate the existence and possible sources of multimodality in the aerodynamic shape optimization of a rectangular wing. Our test case, specified by the ADODG Case 6, has a high dimensionality design space and a large degree of flexibility within that design space. We study several subproblems of this benchmark test case and analyze the multimodality introduced by each set of variables considering both inviscid and viscous analysis. This methodical approach allows us to isolate a few instances where certain parametrizations create a multimodal design space. However, we demonstrate that these occurrences of multimodality are either due to modeling inaccuracies or can be curbed by the application of practical design constraints. Additionally, we find that the shape of the optimized wing is highly dependent on the interplay between induced and parasite drag, providing more incentive to consider viscous effects in the analysis. As an example, when minimizing drag with respect to chord distribution, there are multiple local minima at low C_L , but a single global minimum (the elliptical planform) at high C_L . These results will help other researchers to avoid some of the pitfalls that could lead to multimodality in wing optimization.

1 Introduction

In 1799, Sir George Cayley first conceived of the fixed-wing airplane [1]. Despite its long history, the airplane wing is still the subject of active research and design efforts. In recent years, the environmental impact of air traffic and rising fuel prices have driven continuing improvement in wing and more broadly, aircraft performance. Much has changed since Cayley's time however, and the field of aeronautical engineering has been transformed due to advances in scientific computing. The burden of preliminary design and analysis has largely shifted from the laboratory to the computer thanks to the accessibility and efficiency of CFD. Although CFD is an invaluable tool for aerodynamic analysis, wing design is a highly multidisciplinary undertaking. As such, the development of multidisciplinary design optimization (MDO) has been an equally important addition to the designer's toolkit [2]. MDO is concerned with considering multiple facets of the design problem simultaneously, as opposed to sequentially [3]. This paradigm change in the design process enables the accurate representation of coupling between physical disciplines and the precise balancing of systemic trade-offs to achieve an optimal result.

Many of the applications of MDO to aircraft design have been targeted at optimizing specific configurations or features. In one of the first demonstrations of numerical optimization in wing design, Hicks and Henne [4] modified the design of a swept wing with airfoil shape and twist variables to reduce wave drag and increase L/D . Jameson et al. [5] optimized a wide-body wing and wing-fuselage configuration and recovered a shock-free surface. Several studies have investigated the optimization of winglets [6–8]. Martins et al. [9] accomplished a CFD-based aerostructural optimization of a supersonic business jet. More recently, a number of studies have been published on the high-fidelity design optimization of the Common Research Model (CRM) benchmark wing [10–13]. Application of MDO to future aircraft concepts like the blended wing-body [14] and D8 [15] has also been a fruitful area of research.

¹Ph.D. Candidate, AIAA Student Member.

²Visiting Scholar

³Research Investigator, AIAA Senior Member.

⁴Professor, AIAA Associate Fellow.

These results have strengthened the credibility of MDO and its applicability to practical aircraft design problems. However, these studies focused on the refinement of existing technologies, or making good aircraft better. There are sound reasons for starting an optimization problem from a design that already employs state-of-the-art knowledge and techniques. In high-fidelity aerostructural optimizations for example, starting the optimization with a poor initial design can prevent the optimizer from finding a feasible solution. However, in the long term, we want to make MDO robust enough that it can be used to explore the full design space and discover novel designs. Ideally we would be able to start an optimization from a sphere and recover an airplane that would be optimally suited to its mission requirements. Gagnon and Zingg [16] tested the plausibility of this endeavor and succeeded in transforming a half-sphere into a blended wing-body shape by maximizing L/D . There is a growing body of research that seeks to improve the methods for exploratory wing design optimization. For example, Jansen et al. [8] used a panel aerodynamic model coupled with a beam finite-element model to optimize a nonplanar wing composed of four connected wing panels that could assume any shape. They were able to recover, in order of increasing design constraints the following configurations: box-wing, C-wing, raked wingtip, and winglet. Hicken and Zingg [17] demonstrated the use of a B-spline geometry parametrization to enable the large shape changes expected in an exploratory optimization study. They minimized the induced drag of a generic rectangular wing and produced a wing with increased span and vertical extent. Kenway et al. [18] developed a robust parametrization scheme and mesh warping algorithm that enabled the aerostructural planform optimization of a turbo-prop transport aircraft wing.

One of the difficulties associated with exploratory optimization is the possibility of multimodality in the design space. Generally, gradient-based optimization algorithms are used in aerodynamic shape optimization due to the need to reduce the number of computationally expensive function evaluations. However, gradient-based algorithms converge to a single local optimum, and require multiple starting points to find out if there are multiple local minima. Chernukhin and Zingg [19] conducted a lift-constrained induced drag minimization of a rectangular wing with respect to planform variables. They reported finding seven local minima when starting from 192 random perturbations of the baseline geometry. Two explanations arise for this apparent multimodality. The first, at the suggestion of the authors of that study, is that such local minima do exist in reality, but that design constraints render most of them infeasible in practical design problems. The alternative explanation is that the multimodality is an artifact of modeling inaccuracies or a restrictive parametrization. For example, the study in question was conducted with Euler CFD and it might be the case that considering the effects of viscosity would eliminate the multimodality in the design space. There is evidence that high-fidelity analysis may be the key to attenuating the multimodality in the design space. Lyu et al. [10] were able to recover very similar solutions in the shape optimization of the CRM even when starting from wings with randomly perturbed shapes. They concluded that the design space for this problem was basically unimodal and that the optimizer was able to converge on a confined, flat region near the global optimum in every case. While this result seems like strong evidence against multimodality in the case of local shape variables given sufficient modeling fidelity, we are interested to see if the same holds true for optimization involving large changes in the wing planform shape.

The AIAA Aerodynamic Design Optimization Discussion Group (ADODG) introduced a sixth test case with the purpose of studying multimodality in exploratory wing optimization⁵. This test case differs from the CRM wing optimization in that it gives the optimizer much more freedom to modify the shape of the baseline wing. Practical design constraints are neglected in favor of liberating the optimizer to explore the entire design space and perhaps discover a novel optimal design for the prescribed conditions. We use this test case as a foundation for a series of optimization studies in which we explore the physical reasons for multimodality in aerodynamic wing design. In Section 2, we introduce the test case and the tools used to perform the optimizations. Before reporting the results of the full test case, we investigate in Section 3 the implications of including each design variable with some preliminary optimization studies. Then we analyze the full case and provide some concluding commentaries in Sections 4 and 5, respectively. In this work, our purpose is not necessarily to do an exhaustive search of the design space to find all possible local minima, or even to rate this problem in terms of its multimodality. Other researchers have done excellent work in addressing these goals [20, 21]. Our aim is to show whether or not multiple local minima exist, clarify the factors influencing multimodality in the problem, and perhaps elucidate whether or not multimodality should

⁵<https://info.aiaa.org/tac/ASG/APATC/AeroDesignOpt-DG/default.aspx>

Table 1: Baseline geometry performance at $C_L = 0.2625$

Grid	Cells	α	C_D (cts)	$C_{D,v}$ (cts)	$C_{D,p}$ (cts)
Euler L3	180,992	3.023	42.694		
Euler L2	1,447,936	3.030	38.997		
RANS L2	306,432	3.206	144.013	76.378	67.636
RANS L1	2,451,456	3.205	130.545	69.702	60.843

be a concern for practical applications of MDO in aircraft design.

2 Methodology

2.1 Multi-Fidelity Approach

The main goal of the ADODG Case 6 optimization problem is to study the existence of multiple local minima in the design space. In addition to this primary goal, we seek to understand whether such local minima reflect the real physics involved, or are merely artifacts of the modeling and discretization errors. In addressing these two goals, we found it useful to combine results from multiple sources of information. All told, we use three different physics models to analyze the aerodynamic performance of the wing: the Reynolds-Averaged Navier-Stokes (RANS) equations with a Spalart-Allmaras (SA) turbulence model, the compressible Euler equations, and a vortex-lattice method (VLM). The RANS and Euler equations are solved using ADflow [22, 23] and the VLM is implemented in OpenAeroStruct⁶. The details of these solvers and their respective workflows are described in the following two sections.

2.2 High-fidelity Optimization: MACH

The MDO of aircraft configurations with high fidelity (MACH) framework offers a powerful, automated approach to aircraft design [24]. For aerodynamic shape optimization, MACH provides a hyperbolic mesh generator, a free-form deformation geometry parametrization scheme [18], an unstructured mesh warping module [18], and ADflow, a finite-volume CFD solver for cell-centered multiblock and overset meshes. ADflow solves the compressible Euler, laminar Navier-Stokes, and RANS equations with a second-order accurate spatial discretization. For the Euler-based optimizations conducted in this study, we use the diagonalized alternating direction implicit (DADI) algorithm for the initial multigrid iterations and then switch to a Newton-Krylov (NK) solver to tightly converge the residual for each solution. For the RANS analysis, we use a Runge-Kutta algorithm for the multigrid and the same approach with the NK solver. The greatest benefit of the MACH framework for optimization is that each of the modules embedded in the optimization loop provides efficient, accurate gradient computation in addition to its primary function.

The baseline wing geometry for the Euler and RANS analyses is planar with a chord of 1.0 m and a NACA 0012 airfoil cross-section. The wingtip cap is a perfect revolution about the airfoil chord line and adds 0.06 m to the 3.0 m rectangular portion of the wing, bringing the total semispan to 3.06 m. The Euler geometry has a sharp trailing edge, while the RANS geometry has a blunt trailing edge with a thickness of 2.52 mm. We generate the surface meshes using Ansys ICEM CFD and extrude the volume meshes using hyperbolical marching. The meshes are oriented with the x-axis in the streamwise direction, the z-axis out the wing, and the y-axis in the vertical direction. The quality of these meshes is tested in a grid convergence study at the nominal baseline condition ($M = 0.5$, $Re = 5 \times 10^6$, $C_L = 0.2625$), the results of which are plotted in Fig. 1. The drag values converge nearly linearly as the number of cells to the power of $2/3$ is increased, indicating asymptotic convergence. Table 1 lists the data for the baseline grids that are used in the optimization studies.

We manipulate the grids using free-form deformation (FFD) and an inverse-distance mesh-warping algorithm [18]. The FFD approach requires the user to generate a structured volume mesh that encompasses the geometry, represented in Fig. 2 by the 3D grid of black points and lines. The control points of this FFD volume are used to create a trivariate B-spline mapping of its interior. This mapping furnishes the parametric position of each node of the embedded surface mesh. As the positions of the FFD control points change, the embedded geometry deforms continuously according to the B-spline mapping. The design variables are set

⁶<https://github.com/mdolab/OpenAeroStruct>

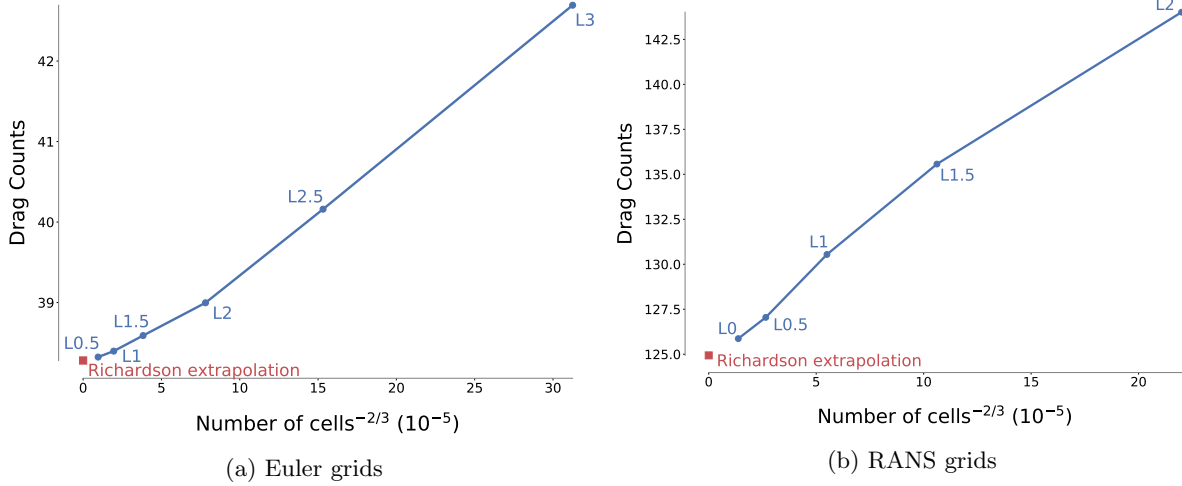


Figure 1: Grid convergence study at $M = 0.5$, $Re = 5 \times 10^6$ and $C_L = 0.2625$.

up to manipulate the position of the control points in order to enact local or global changes to the shape of the embedded geometry. Deformations of the surface mesh are then propagated out to the volume mesh by way of the mesh-warping algorithm. The derivatives of the volume mesh nodes with respect to the surface mesh nodes are computed using automatic differentiation (AD). The derivatives of the surface mesh nodes with respect to the FFD control points are computed analytically and the derivatives of the FFD control points with respect to the design variables are computed using the complex-step method.

Our parametrization uses both global and local design variables. The global design variables act on a group of the FFD control points, facilitating large-scale deformations. Thus, a given global design variable will produce a nonzero derivative for multiple control points. The local shape design variables code for the individual displacement of each control point in a user-specified direction. It follows that the jacobian of the control points with respect to the local design variables is the identity matrix. When multiple design variables affect a given control point, the operations are combined linearly. Fig. 2 depicts the FFD volume with the control points as black dots and the design variables definitions shown with color-coordinated arrows. The nominal FFD volume has nine spanwise control sections and 12 chordwise control points per section with half of the points on the upper surface and the other half on the lower surface. The global variables are linked to the displacement and rotation of axial control points (shown as red squares) along the reference axis. Each of these axial control points dictates the global movement of an entire FFD control section. For example, the blended winglet shown in Fig. 2 was created by setting the dihedral variable with the following values:

$$x_{\text{dihedral}} = [0, 0, 0, 0, 0.03, 0.09, 0.24, 0.51, 1]$$

The user defined function for the dihedral variable displaces each of the nine axial control points in the y -direction according to the corresponding value in x_{dihedral} . Additionally, the dihedral function is set up to rotate the FFD sections so that they remain perpendicular to the reference axis. The local shape variables are handled differently. Each spanwise control section is assigned a unique reference frame with the section plane normal as the \hat{e}_k axis, the \hat{e}_i axis aligned with the streamwise direction, and $\hat{e}_j = \hat{e}_k \times \hat{e}_i$. The local shape variables control the movement of the control points along their respective \hat{e}_j axes. When the reference axis is displaced vertically, the control sections and their respective reference frames are automatically rotated to remain perpendicular to it. This behavior is depicted in the formation of the blended winglet on the wing in Fig. 2. The displacement vectors of the local shape variables in the wingtip control section are rotated as the winglet forms to allow sectional control of the airfoil section of the winglet. This functionality ensures that the wing surface does not shear, causing negative volumes, when large changes in dihedral are introduced.




Figure 2: Wing embedded in the FFD volume with design variable definitions.

2.3 Low-fidelity Optimization: OpenAeroStruct

OpenAeroStruct [25] is an open-source low-fidelity aerostructural optimization suite developed using the OpenMDAO framework [26]. The aerodynamic analysis in OpenAeroStruct is performed using a vortex lattice method (VLM) to compute induced drag and a modified flat-plate skin-friction drag approximation to estimate parasite drag. These low-fidelity models provide reasonable estimates at a low computational cost. A single analysis takes less than one second and a full optimization takes on the order of 10 seconds on a single processor. For the analyses in this study, the baseline geometry consists of a $1\text{ m} \times 3.06\text{ m}$ rectangular half-wing discretized by 50 spanwise panels and mirrored across the symmetry plane. The thickness-to-chord ratio and location of maximum thickness from the NACA 0012 airfoil are used in the computation of skin friction drag. The geometry is parametrized using B-splines to interpolate variable changes to the geometry. Using B-splines allows a reduction in the number of design variables so that 50 panels can be manipulated with only 9 spanwise control points.

2.4 Optimizer

Both MACH and OpenAeroStruct are optimizer-independent, but for this study we use SNOPT [27] exclusively. SNOPT is a sequential quadratic programming, gradient-based optimizer which has been vetted for aerodynamic shape optimization and aerostructural design optimization [10–12, 28]. All of the optimization results presented herein are converged to a feasibility tolerance of at least 1×10^{-6} , which corresponds to six digits of accuracy in the lift coefficient. The optimality tolerance for every result is at least 2.5×10^{-4} , and for the majority of the results, it is less than 1×10^{-5} .

2.5 Optimization Problem

The complete optimization problem of ADODG Case 6 (with some minor alterations) is defined in Table 2. Twist variables are defined at every axial control point except at the root and α is used to match the C_L constraint. The chord variables affect all nine axial control points. In the Euler cases, when local shape variables are inactive, the chord scales while maintaining constant t/c . Dihedral is defined as the vertical displacement of each axial control point. Additionally, the FFD section corresponding to each axial control point rotates to align the twist rotation and any shape variable displacements to be perpendicular to the wing surface. Sweep is defined as the streamwise displacement of each axial control point. In the official ADODG case, the limits on sweep extend to 1m forward and backward in the streamwise direction. In our optimizations, we reduced these limits to 0.5m to avoid problems with excessive mesh warping. Both dihedral and sweep are fixed at the root. As explained previously, the local shape variables perturb the wing cross-section perpendicular to the wing surface, such that they have some dependence on the dihedral variable. All control points in a given section are perturbed in a uniform direction, as indicated in Fig. 2.

Table 2: ADODG Case 6 Optimization Problem Statement

Category	Name	Quantity	Lower	Upper	Units
Objective	C_D	1	–	–	–
Variables	α	1	–3.0	6.0	degrees
	γ	8	–3.12	3.12	degrees
	c	9	0.45	1.55	m
	dihedral	8	–0.45	0.45	m
	b	1	2.46	3.67	m
	sweep	8	–0.5	0.5	m
	shape	108	–0.5	0.5	m
	Total	143			
Constraints	C_L	1	0.2625	0.2625	–
	S	1	3.06	3.06	m ²
	V	1	V_0	–	m ³
	t	108	$0.5t_0$	$1.5t_0$	m
	Total	112			

The planform area, S , is computed as the area of the wing projected onto the x-z plane. Constraints for volume (V) and thickness (t) are handled by first setting up a 2D grid of points inside the surface of the wing. Then these points are projected to the surface of the wing to create a 3D grid confined within the wing. We compute V as the sum of the cell volumes and t as the difference between the projected points on the upper and lower surface. The thickness constraints are evaluated at ten uniformly-spaced chord-wise locations ranging from $0.005c$ to $0.99c$ for ten sections along the span. There are an additional eight thickness constraints added in the wingtip cap, making a total of 108 thickness constraints. We removed the root bending moment constraint from the official optimization problem. The root bending moment constraint is related to structural requirements, and in these results we want to consider solely the influence of aerodynamics on the wing design. The optimization cases treated in Section 3 are subproblems of this full problem, and the variables and constraints are defined as stipulated in this full problem description unless otherwise stated. All optimizations are run at Mach 0.5 and a Reynolds number of 5 million.

3 Parametrization Studies

3.1 Twist Optimization

We begin with a simple twist optimization problem as a means of verification for our optimization framework. The twist optimization case has a long theoretical history and has also been extensively studied as a numerical optimization problem in ADODG Case 3 [29–31]. The elliptical twist distribution is well known as the theoretical optimum for this case because it generates the constant spanwise downwash required to minimize induced drag [32]. This theoretical result is a useful metric with which to gauge the performance of our optimization framework. Fig. 3 shows the optimal twist and lift distributions for the three levels of fidelity. We see very similar trends from each of the analyses, but there is a noticeable offset between the lift distributions from ADflow and the VLM results. This discrepancy is due to the rounded wingtip cap used for the Euler and RANS geometries. In the lifting-line model, the entire span is used to generate lift, whereas with the wingtip cap, the leading and trailing edges are truncated at 3 m of span and the last 0.06 m of span is incapable of generating the lift required to complete the elliptical distribution. As a result, the optimizer converges to a wing that generates an elliptical lift distribution extending from the root to the edge of the wingtip cap, represented by the shifted elliptical curve in Fig. 3. The VLM results match the original elliptical curve because they are obtained from a 2D surface for which the problems discussed above are irrelevant. Table 3 lists the drag counts of the optimized wings and the percent difference from the baseline drag value, $\% \Delta C_D$. As an added verification, we experimented with varying the number and spacing of the twist variables along the span and also started the optimization from ten random starting points. All of these variations yielded consistent results, which leads us to confirm the theoretical assertion that there is a single twist distribution that produces the lowest drag.

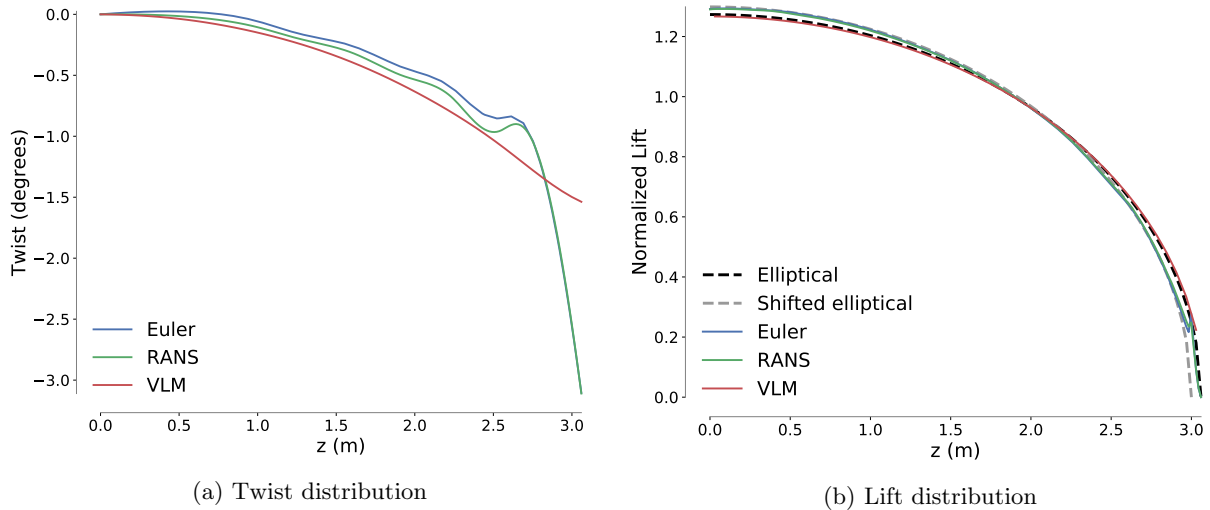


Figure 3: Wings optimized with respect to twist variables generate an elliptical lift distribution.

Table 3: Twist optimization results

Case	Grid	C_L	Drag counts	$\% \Delta C_D$
Euler	L3	0.2625	42.105	-1.38
RANS	L2	0.2625	143.377	-0.44
VLM	—	0.2625	128.779	-0.42

3.2 Chord Optimization

Different behaviors arise when drag is minimized by varying the chord distribution while keeping the twist constant. Theoretically, the elliptical chord distribution should be optimal for minimizing induced drag, and historically, this concept has been put to the test in the design of actual aircraft, most notably the Supermarine Spitfire. The optimization problem is to minimize drag with respect to the chord distribution, subject to the constraint that $C_L = 0.2625$. Since the lift coefficient is normalized by S , which may vary with changes in the chord, we must also constrain S to be constant. The inviscid and viscous results are henceforth discussed separately for each case to better explore and highlight the unique characteristics of each.

3.2.1 Euler

For the Euler chord optimizations we use the L3 mesh and change the bounds on the chord variables to (0.1, 2.0) to allow the planform to match an ellipse as close as possible. As in the twist optimization, the inviscid chord optimization yields predictable results. We take this predictability as an opportunity to test the sensitivity of the result to the chosen parameterization. We vary the number of spanwise FFD sections and corresponding chord variables and also compare the difference between scaling about the trailing edge and the quarter-chord. As shown in Fig. 4, regardless of these modifications, each optimization converges to an elliptical lift distribution. The effect of varying the number of FFD control sections is marginal. However, we can see from the results in Table 4 that scaling the chord about the trailing edge allows more than double the drag reduction compared with scaling about the quarter-chord. This discrepancy was previously observed in a thorough investigation of the differences in induced drag performance between elliptical and crescent (elliptical with straight trailing edge) planforms by Smith and Kroo [33]. Theoretically, both of these planforms should produce a constant downwash and generate minimum induced drag. Smith and Kroo point out that these optima were derived using lifting-line theory, which relies on a tip vorticity of infinite strength to complete the constant downwash distribution. In nature, such a singularity is impossible, and therefore, even for a perfectly elliptical wing, there is a slight upwash at the tip. They found that the crescent




Figure 4: Elliptical planform and lift distributions obtained from optimization with respect to chord variables.

Table 4: Chord optimization results

Case	Grid	C_L	Drag Counts	$\% \Delta C_D$
Euler 1/4 chord	L3	0.2625	41.726	-2.27
Euler 6 trailing edge	L3	0.2625	40.439	-5.28
Euler 9 trailing edge	L3	0.2625	40.412	-5.35
Euler 17 trailing edge	L3	0.2625	40.356	-5.48
RANS	L2	0.2625	142.544	-1.02
RANS mode 1	L1	0.2625	129.214	-1.02
RANS mode 2	L1	0.2625	129.178	-1.05
RANS monotonic	L1	0.2625	129.296	-0.96
RANS	L1	0.5	234.663	-1.23
RANS	L1	0.8	503.049	-3.85

wing produced a more constant downwash near the tip, which led to a more nearly elliptical span loading and a lower overall induced drag than the elliptical wing. This finding may seem to contradict Munk’s well-known Stagger Theorem, which maintains that any streamwise arrangement of lifting elements should generate the same induced drag [32]. However, the Stagger Theorem is also based on lifting-line theory, and suffers from the same assumptions in the translation to reality. Thus, although theoretically, a straight wing and a swept wing of the same span and chord distribution generate equal induced drag, in reality, differences in tip downwash most likely cause slight differences in induced drag too.

To test the multimodality of this problem, we started the optimization from ten random starting points. Nine of the starting points converged to the elliptical planform and one failed prematurely due to mesh warping errors. These results indicate that there is no multimodality in chord optimization for inviscid flow.

3.2.2 RANS

When adding viscous effects to the chord optimization, it is important to consider the trade-off between induced drag and parasite drag, and its relationship to chord length. For the purposes of this discussion, since we are dealing with subsonic flow of a fairly streamlined geometry, we can reasonably assume that the majority of the viscous contribution to drag is made up of skin friction drag. However we do recognize that viscosity also introduces pressure drag due to separation, especially as C_L increases. While induced drag is sensitive to the spanwise distribution of lift, skin friction drag is highly dependent on the local chord length. The shear stress at the wall is directly related to the velocity gradient normal to the wall. For 2D laminar




Figure 5: The viscous drag on a flat plate in laminar flow can be reduced by increasing the average BL thickness over the plate. ($Re = 10^6$)

flow over a flat plate,

$$\tau_w = \mu \frac{du}{dy} \quad (1)$$

This equation provides an approximation to the shear stress on an airfoil. At the leading edge, the boundary layer is very small and the velocity changes rapidly over a small distance, resulting in a large shear stress. As x/c increases, the boundary layer (BL) fills out and the velocity gradient at the wall becomes much more mild. The result is that extending the chord reduces the average drag per unit length. This relationship is well known and Blasius [34] provided the following analytic solution for the flat plate case,

$$c_d = \frac{1.328}{\sqrt{Re}}, \quad (2)$$

where the Reynolds number is based on the local chord. Multiplying by the local chord length we find that for constant flow conditions, the skin friction drag per unit span is proportional to the square root of the local chord length, i.e.,

$$d \propto \sqrt{c} \quad (3)$$

To illustrate this point, imagine we want to minimize the drag of a flat plate in laminar flow at zero angle of attack with a fixed planform area. If the chord distribution and span are variables, the chord distribution will grow to its upper limit and the span will adjust to satisfy the area constraint. However, if span is fixed, or has a hard lower limit, an interesting compromise takes place. The optimal chord distribution will have the maximum possible extent of the span at the upper limit of the chord variable. For example, if we split our plate into two independent sections and incrementally add Δc to the chord of one section while subtracting the same Δc from the other, we get a decrease in skin friction drag as shown in Fig.5. In the absence of other constraints, lower skin friction drag can always be achieved by transplanting wing area from a thin-BL region to a thick-BL region. The spanwise location of the maximum chord region is irrelevant, and as such, a purely skin friction drag minimization problem theoretically has an infinite number of local minima.

The problem becomes more complex when the objective function is a combination of both skin friction and induced drag. Minimum skin friction drag favors radical changes in the chord distribution to maximize thick-BL coverage, but optimal inviscid drag calls for an elliptically tapered wing. The optimal planform shape will balance these considerations taking into account the relative weights of each drag component. For instance, if drag is mostly induced (i.e., at high C_L), the planform will take on a nearly elliptical profile. However, if we make some slight modifications to the planform, we may get some improvement in the skin friction drag while not straying too far from the elliptical lift distribution. This line of reasoning helps to explain the results we get from the RANS chord optimization, shown in Fig.6. These results are obtained

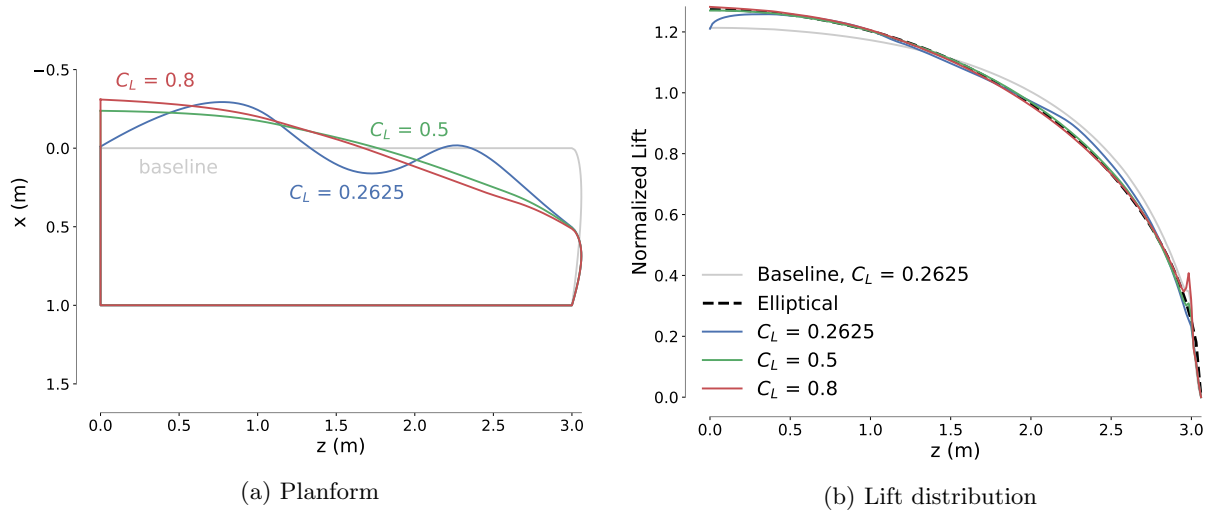


Figure 6: The optimal chord distribution for viscous flow is heavily dependent on C_L .

using the RANS L1 mesh. For a given C_L , we can estimate the minimum possible contribution of induced drag to the total drag using Prandtl's induced drag approximation $C_{D,i} = C_L/\pi AR$. For $C_L = 0.2625$, parasite drag makes up roughly 70% of the total drag, and the optimal planform shape is far from elliptical. As C_L increases and induced drag increases relative to parasite drag, the optimal design more closely approaches the elliptical planform. However, even when C_L is relatively low, the optimal lift distribution oscillates close to the elliptical profile. This is not surprising because in reality, even a rectangular wing has a span loading not far removed from the elliptical ideal. Fig. 6 shows the lift distribution of the baseline wing at $C_L = 0.2625$. As seen in the Euler results in Table 4, going from a rectangular wing to an elliptical wing produces a 2–6% drag reduction. So, while the elliptical planform is ideal, the rectangular wing is certainly not a terrible starting point. This accounts for the rather meager gains in aerodynamic performance observed in the twist and chord optimizations overall.

To address whether this trade-off between parasite and induced drag leads to multimodality at low C_L values, we first turn to a simpler aerodynamics model. Fig. 7 shows the results of optimizing 20 random initial planforms at five different C_L values with OpenAeroStruct. A trade-off exists between minimizing skin friction drag by maximizing chord and forming an elliptic planform to minimize induced drag. As expected, based on our previous discussion, the deviations from the elliptical lift distribution increase dramatically as C_L drops to zero. A corresponding decrease in skin friction drag is small but noticeable. As C_L increases, induced drag makes up a greater portion of the total drag, and thus the planform approaches an elliptical shape and multimodality decreases. In the high-fidelity case we see similar trends. We optimize three randomly generated planforms at $C_L = 0.2625$ with the RANS L1 mesh and the optimizer converges to two different planform modes. Wary of this result, we vary the number of control points used as design variables to rule out the possibility that our parametrization is biasing the results. We optimize the baseline geometry with 6 and 17 spanwise variables and compare the result with the baseline optimization with 9 spanwise variables. Despite the variance in flexibility allowed by the number of spanwise control points, the optimizer converges on a very similar, albeit not identical, planform shape in all three cases. These experiments suggest that there are a limited number of modes with which the optimizer can minimize skin friction drag while still maintaining a sufficiently elliptical lift distribution so that the total drag decreases. The results for both of these tests are shown in Fig. 8. Incidentally, we began the chord optimization tests with the RANS L2 mesh, but found a lack of multimodality in the low C_L cases. Since this did not agree with our hypothesis, we refined the mesh and found that the finer mesh yielded sharper spanwise curvature due to the refinement of the spanwise grid spacing. We surmise that the coarseness of the L2 mesh causes an increase in drag when the optimizer attempts to create the large-amplitude spanwise variations seen in Fig. 8 and thus artificially limits the design space. However, it should be noted that the same physical phenomenon is apparent in both meshes, although its effect is dampened due to the coarseness of the mesh.

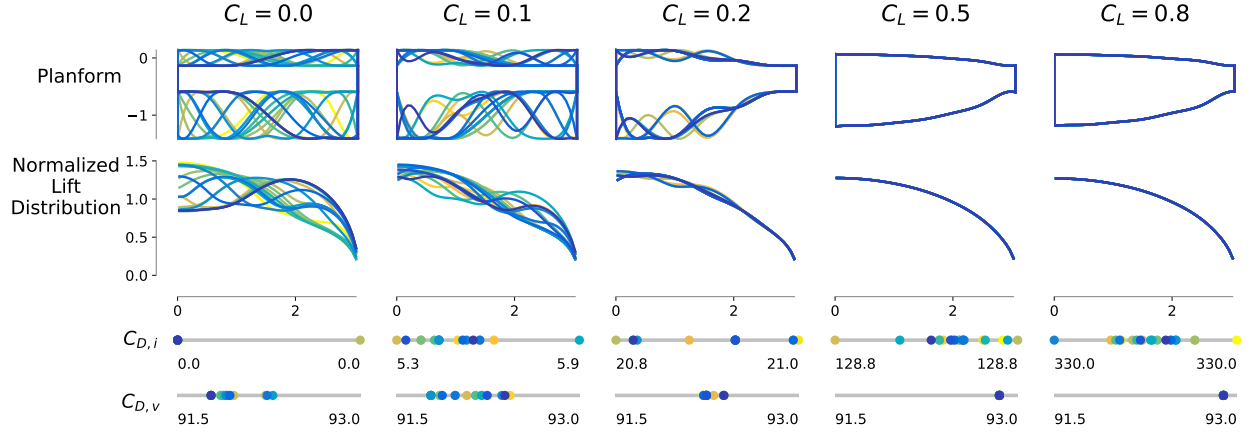


Figure 7: Optimization using OpenAeroStruct confirms multiple optimal chord distributions at low C_L .

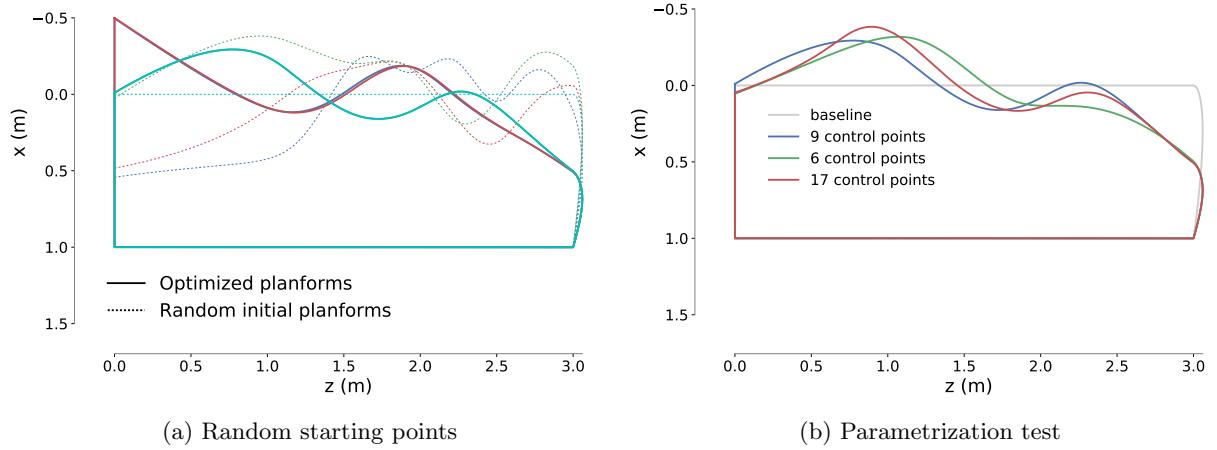


Figure 8: (a) Three optimizations from random chord distributions yield two local minima. (b) Varying the number chord variables changes the optimal distribution slightly.

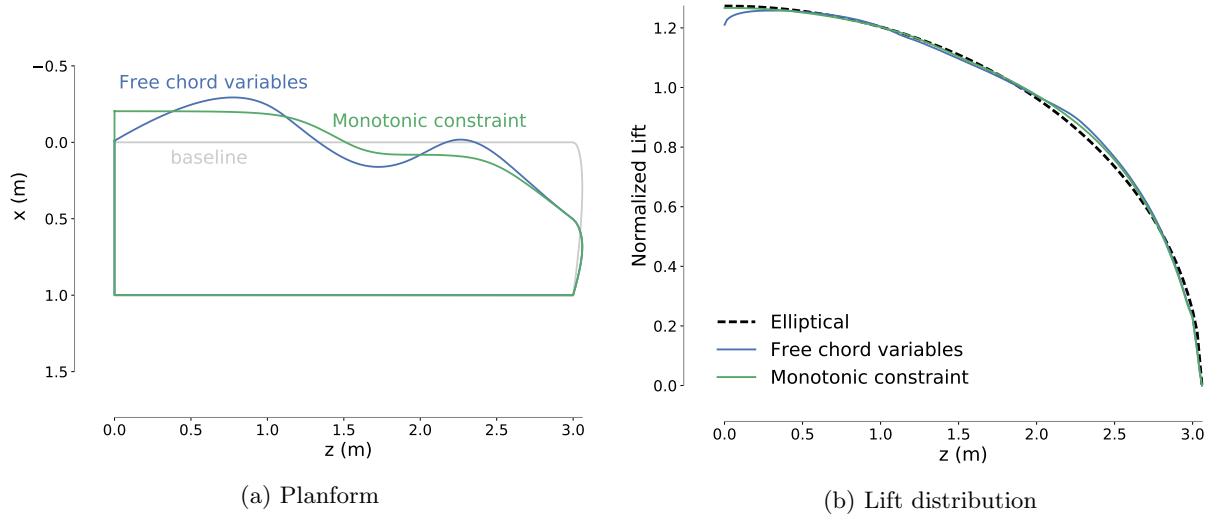


Figure 9: Forcing the chord distribution to decrease monotonically increases drag by less than one count.

While these results are intriguing, we are ultimately interested to see how much of a benefit these spanwise oscillations really provide. The VLM results suggest a drag decrease on the order of one to two counts. We add a monotonic constraint to the optimization of the L1 grid as a crude substitute for a more meaningful constraint, such as manufacturing cost and structural considerations. This constraint forces the chord distribution to decrease monotonically from root to tip. The result is plotted in Fig.9 and detailed in Table 4. Adding the monotonic constraint only increases the drag by a fraction of a drag count. This difference is not meaningful because it is within the modeling error. Furthermore, even if this difference were meaningful, it would be hard to justify the added manufacturing costs and structural penalties that would accompany building a wing with such spanwise curvature.

3.3 Chord and Twist Optimization

Since we can reach an elliptical distribution with only chord design variables or only twist design variables, one may assume that a combination of both twist and chord variables would be redundant and lead to multimodality. We explore this possibility in this section.

3.3.1 Euler

For the chord and twist optimization with Euler analysis we use the Euler L3 mesh and nine axial control points, corresponding to nine chord variables and eight twist variables. Once again, the limits on the chord variables are relaxed to $0.1 \leq c \leq 2.0$ for the Euler cases. When initialized with ten different geometries, each generated with random chord and twist distributions, the optimizer converges to a single optimal solution. The optimal chord and twist distributions are superimposed over the randomly generated seeds in Fig. 10. The uniformity of these results suggest that the chord and twist variables are not necessarily interchangeable. In Table 5, we see that the addition of twist variables makes a negligible improvement in drag compared with the chord-only optimization (Table 4).

3.3.2 RANS

The RANS optimizations are run using the RANS L1 mesh and the nominal 9-section FFD. The results are displayed in Fig.11. We expect to see the oscillatory behavior observed at low C_L in the chord optimization to be exaggerated in this case because the twist can compensate for the deviations from the elliptical lift distribution that oscillations in the chord distribution would otherwise cause. The higher C_L results, like the Euler results, show little evidence of multimodality due to redundancy in the variables. For $C_L = 0.2625$, the oscillations in the chord distribution are more prominent, which implies that the addition of twist variables grants more freedom to the chord variables to minimize skin friction drag. Once again, the improvement in




Figure 10: Optimizing from 10 unique starting points with respect to chord and twist variables using Euler analysis yields a single optimum.

Table 5: Chord and twist optimization results

Case	Grid	C_L	Drag Counts	$\% \Delta C_D$
Euler 1/4 chord	L3	0.2625	41.517	-2.76
Euler	L3	0.2625	40.408	-5.35
RANS	L1	0.2625	128.802	-1.34
RANS	L1	0.5	233.591	-1.68
RANS	L1	0.8	480.496	-8.16

drag gained from the chord and twist optimization is fairly insignificant due to the fact that there is only so much that can be improved with twist and chord variables.

3.4 Dihedral and Twist Optimization

Nonplanar wings have the potential to reduce induced drag beyond what is attainable with a planar wing and an elliptical lift distribution. Previous optimization studies have verified this result [8, 35]. We include twist variables in this subproblem because we want to allow the optimizer to converge to an elliptically loaded planar wing if that is the optimal design. However, we first assess the parametrization of the dihedral variables by conducting an optimization with just dihedral. Since we expect the most variation in dihedral to occur toward the wingtip, we want to make sure that our parametrization allows enough flexibility to capture the optimal shape. We vary the number of control points and experiment with spacing them uniformly along the span and clustering them more heavily toward the wingtip. In Fig. 12 we compare the winglet-down optimized results for six different possible parametrizations. All six parametrizations achieved an elliptical lift distribution, but the cases with only six variables prevented the optimizer from converging to the optimal winglet cant angle. We choose to use the FFD with nine control points clustered toward the wingtip as the nominal FFD because it gives a sufficient degree of freedom to the optimizer for the kind of design space exploration we seek to do. Note that we recover two local minima in the dihedral-only optimizations: an upturned winglet and a downturned winglet. The upturned winglet shapes are achieved when starting the optimization from random starting points; Fig.12 shows only the optima found when starting from the baseline geometry.

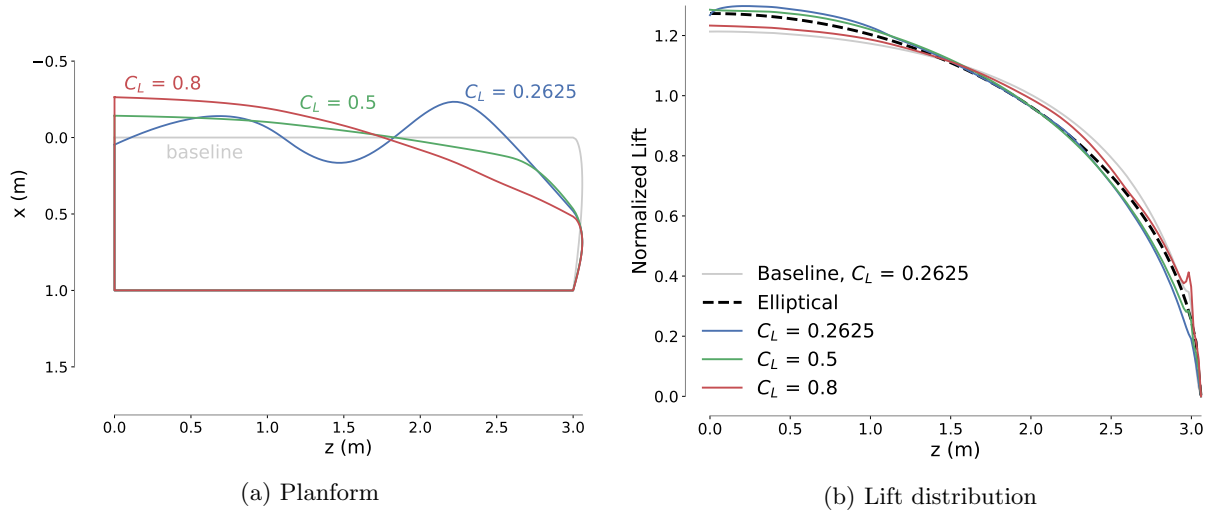


Figure 11: An optimization with both twist and chord variables tolerates more variation in the chord distribution at a lower C_L . As C_L increases, the planform shape oscillations disappear.

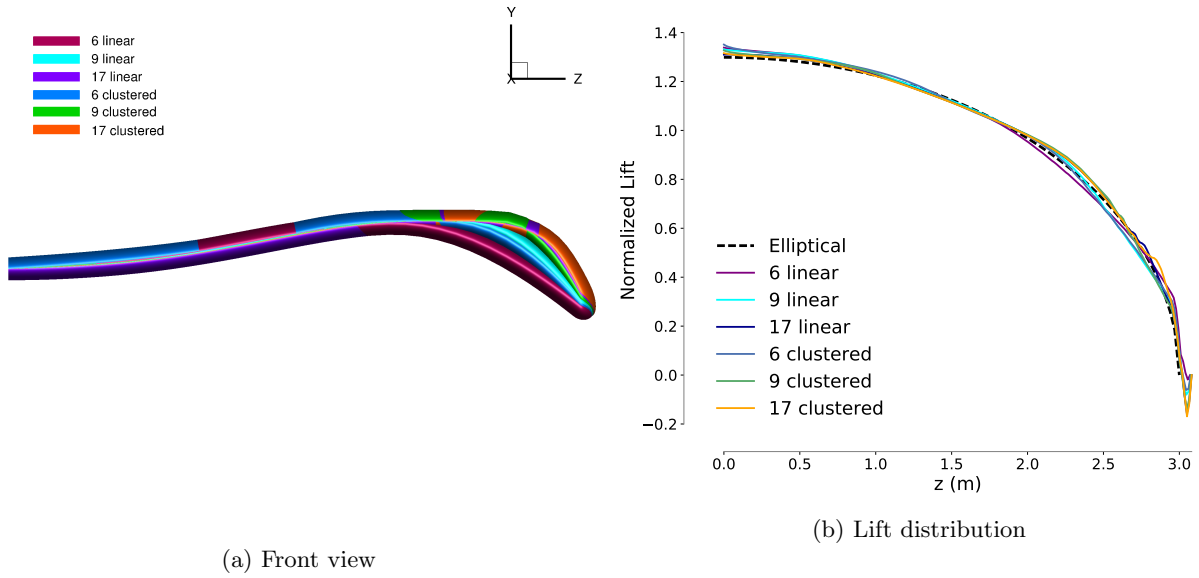


Figure 12: Comparison of different quantities and spanwise distributions of control points for an optimization with dihedral variables.

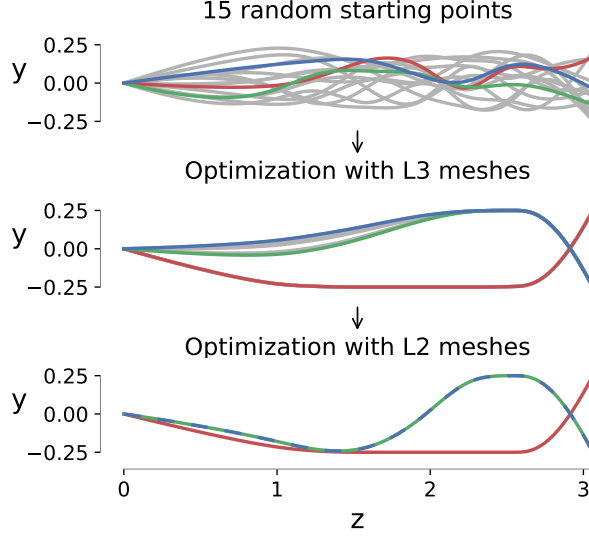


Figure 13: Results of optimizing the dihedral distribution first with L3 meshes and then with L2 meshes.

Table 6: Dihedral and twist optimization results. The $C_L = 0.8$ result was obtained by regenerating the volume mesh to fix problems with the formation of negative volumes.

Case	Grid	C_L	Drag Counts	$\% \Delta C_D$
Euler	L3	0.2625	39.909	-6.52
RANS	L2	0.2625	143.327	-0.48
RANS	L2	0.5	244.969	-2.79
RANS	L2	0.8	455.178	-15.21

3.4.1 Euler

Now we combine twist and dihedral to see if the coupling between these two sets of variables affects the number of optimal designs. We start with running 15 optimizations starting from unique dihedral distributions. These initial distributions can be seen in the top frame of Fig. 13. In the optimization with L3 meshes, the optimizer converges on multiple local minima. Five of the optimizations converge to an upturned winglet shape and the rest converge to a downturned winglet. Those that converge to the winglet-down shape share a very similar wingtip design, but vary in the vertical displacement of the main wing. These differences can be seen in the middle frame of Fig. 13. The next set of optimizations starts from the colored distributions in the middle frame of Fig. 13, but uses L2 meshes. The designs obtained in the optimization with finer meshes are shown in the final frame of Fig. 13. The transition to finer meshes reveals that some of the multimodality observed in the first optimization was due to the coarseness of the mesh, and not to any physical phenomenon. In the end, the optimizer converged to a wing with an upturned winglet and a wing with a downturned winglet. In inviscid flow, increasing the horizontal and vertical extent of the wing helps to decrease drag. In the downturned winglet design, the optimizer has maximized the utilization of vertical space by changing curvature in the middle of the wing. Although this is optimal for inviscid analysis, the penalty in parasite drag due to the increased surface area would make this wing suboptimal in realistic analysis.

3.4.2 RANS

The addition of viscosity to the model activates a trade-off between the induced drag improvement from winglet formation and the rise in parasite drag due to an increase in wetted area. Before running the RANS analyses, we explore the design space using OpenAeroStruct to better understand the implications of this trade-off. Fig. 14 compares the optimization results for three general cases: inviscid analysis, viscous analysis, and viscous analysis with a monotonic constraint. The monotonic constraint forces the optimizer to form a

downturned winglet, but the results are of interest because they give a metric with which to compare the other results. The key in the upper left corner provides labels for the different colors. The three cases are optimized with 5 different starting points, displayed in the left-most column, and 5 different lift conditions ranging from $C_L = 0.1$ to $C_L = 0.8$. The front views of each optimized wing profile are shown in each cell of the grid. Below and to the right of each wing profile, slider bars show the relative drag differences between the viscous-optimized results and the inviscid-optimized result. The horizontal bar shows the percent difference in inviscid drag, while the vertical bar indicates the percent difference in skin friction drag compared with a viscous analysis of the inviscid-optimized result. Moving from the left to the right of the grid, the trend is an increase in nonplanarity for the viscous results. At low C_L values, when the skin friction drag dominates, the optimizer has an incentive to reduce the arc length of the front view and thus reduce the effective wetted area. Since only small changes to the dihedral variables are possible without increasing the wetted area, the optimizer converges to the same optimum regardless of the starting point. As C_L increases, this incentive diminishes and the optimizer tends toward more nonplanar wing shapes where the induced drag can be minimized. The increased freedom to vary the wing dihedral opens up the design space to multiple local minima. The inviscid optimization converges to three local minima and the viscous optimization for $C_L = 0.8$ converges to two optima: an upturned and a downturned winglet. The downturned winglet appears to offer better drag performance compared with the upturned winglet.

Taking the information from OpenAeroStruct as a reference, we optimize the RANS L2 baseline planform with respect to eight dihedral and twist variables for $C_L = 0.2625$, 0.5 , and 0.8 . The results once again exhibit a strong dependence on C_L . As seen in left-hand side of Fig. 15, the low C_L case deflects only slightly from the baseline geometry. As C_L increases, the vertical extent of the wing expands. Interestingly, when chord variables are added into the optimization, the trends seen in the separate chord and dihedral optimizations seem to be linearly combined with no recognizable coupling between dihedral and chord. This combined dihedral, twist, and chord optimization is presented on the right-hand side of Fig. 15. In both of these cases, the $C_L = 0.5$ result converges to a down-turned winglet, while the $C_L = 0.8$ result converges to an upturned winglet. We run a dihedral and twist optimization for three wings with randomly distributed dihedral and twist at $C_L = 0.5$ to investigate the possibility of multiple local minima (Fig. 16). Surprisingly, the optimizer converges to an upturned winglet for all three of these randomly generated wings, despite the fact that a comparison of drag values reveals a preference for the downturned winglet. We conclude that in an optimization where the parametrization allows for winglet formation, a gradient-based optimizer can converge to either an upturned or downturned winglet, depending on the starting position. If removal of one of these optima was desired, it would be trivial to apply a constraint which would force the winglet in the preferred direction.

3.5 Adding span and sweep variables

As a final step before considering the full case, we investigate the effects of adding span and sweep variables to the RANS optimization. Fig. 17 shows the results of two distinct optimization problems. The left-hand side shows a comparison between optimizing with respect to twist, chord, dihedral, and span at $C_L = 0.2625$ and $C_L = 0.5$. The only geometric constraint is planform area. For $C_L = 0.2625$, the optimizer finds greater benefit in maximizing the chord and does not increase span to reach the upper bound. The span does increase from the baseline case though, and the optimizer adds a slight anhedral to the wing. On the other hand, for $C_L = 0.5$, the span reaches the upper bound and an upturned winglet is formed.

When sweep is added to the optimization, we find the optimizer tends to sweep the wingtip back sharply, creating a raked wingtip. When the full bounds recommended in the ADODG case description are used, this sharply swept wing tends to cause mesh warping errors, resulting in negative volumes. To avoid this problem, we reduced the bounds to allow sweep to vary from -0.25 m to 0.25 m. The right-hand side of Fig. 17 shows the results of this optimization. The most noticeable feature of these results is the formation of the raked wingtip. Once again we notice that the optimizer does not extend to the full bounds of the span variable for the lower lift case. This time, for $C_L = 0.5$, the optimizer converges on a downturned winglet. The appearance of the same swept-back wingtip in both local minima, despite the differences in chord, dihedral, and sweep distribution on the main part of the wing, suggests that the sweep at the wingtip is highly beneficial to drag minimization. For a wing at incidence to the freestream tip sweep can function similarly to dihedral, dispersing the shed vorticity over a larger region and reducing vortex drag. Additionally, the wing could be benefiting from more constant downwash due to the tip sweep, as explained

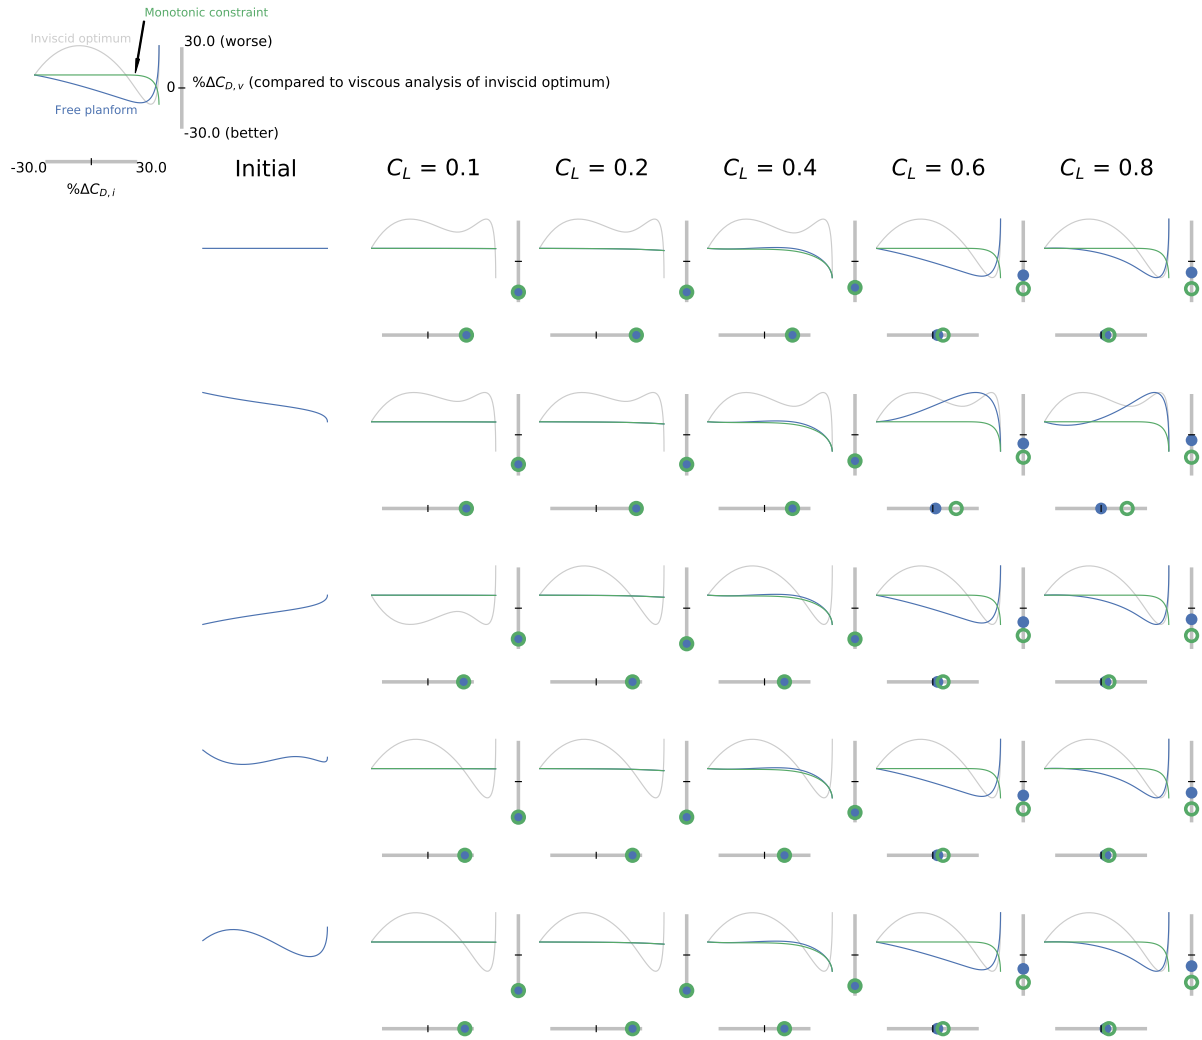



Figure 14: We use OpenAeroStruct to quickly explore the design space of the dihedral and twist optimization.



(a) Dihedral and twist

(b) Dihedral, twist, and chord

Figure 15: (a) A nonplanar wing forms as C_L increases. (b) The chord varies at low C_L and the winglet forms at a high C_L .




Figure 16: The optimizer converges on two local minima for $C_L = 0.5$.




Figure 17: RANS optimization results starting from the baseline rectangular wing at different C_L values.

in Section 3.2.1 regarding the crescent wing.

4 Full Case Optimization

The results of Section 3 emphasized the importance of considering viscous effects when spanwise chord and dihedral variables are included in the optimization problem. Therefore, although the ADODG case calls for Euler analysis, we only report optimization results obtained using RANS analysis in this section. Euler analysis likely would produce multiple local minima, however, our purpose is to study physically meaningful multimodality, and thus we restrict ourselves to RANS analysis. All of these results are produced using the RANS L2 mesh.

The first result presented here is a solution of the optimization problem in Table 2. We run the optimization starting from the baseline geometry and two randomly perturbed starting positions. As seen in Fig.18, the optimizer converges on a different local minimum in each case. Despite the apparent differences in planform shape, the final drag values are remarkably similar (see Case 1 in Table 7). Following the trend of these results, we suppose that starting the optimization from additional points in the design space would reveal more local minima. In keeping with our purpose, we do not attempt such an investigation, but we would like to point out some of the trends that have reappeared from Section 3 in these results. True to our previous hypothesis, the chord variation is multimodal—even more so than before because for a given chord distribution the optimizer has more design variables with which to tailor the wing shape to minimize induced drag. Each optimal design has a unique sweep distribution, but there is a commonality of sharp sweep-back at the wingtip. As expected from our study of the dihedral variables, there is very little y-displacement in any of the three solutions due to the low lift coefficient and the relative unimportance of minimizing induced drag. Instead of pushing the span to its maximum allowable extent, the optimizer uses the available planform area to maximize the chord and reduce skin friction drag. In all three cases, the span increases to a different value, implying that the optimum aspect ratio varies depending on the shape of the wing. In terms of drag reduction, all three solutions improve by nearly 11%. The breakdown between viscous and pressure drag in Table 7 reveals a significant reduction in pressure drag. Up to this point in our investigation we have not used shape variables to tailor the airfoil cross-section of the wing. It appears that the optimizer has taken advantage of the shape variables to streamline the wing shape and reduce drag due to viscous separation. This can also be seen in the pressure plots on the right of Fig. 18, where the flattened profiles differ significantly from the well-known steep pressure gradient of the baseline NACA 0012 airfoils.

In the remaining results, we consider different methods of constraining the problem to manage the multimodality. We first experiment with adding the root bending moment constraint back into the problem. The result is a planform similar to Optimum 1 of in Fig.18, only with a shorter span to satisfy the added constraint and a corresponding increase in drag (see Case 2 in Table 7). Since adding this constraint does not appear to attenuate the multimodality in the chord distribution we do not investigate further. Next we reintroduce the constraint forcing the chord distribution to decrease monotonically from the root to the tip. With this constraint, the optimizer converged on two different local minima in three different runs. The results are shown in Fig.19 and tabulated in Table 7 (Case 3). For both of these optima, the monotonically decreasing chord distribution resulted in slightly better performance than any of the optima found in the previous optimization. Adding the monotonic constraint gave the optimizer freedom to extend the span to its maximum allowable extent, resulting in better induced drag performance. This can be seen in the slightly

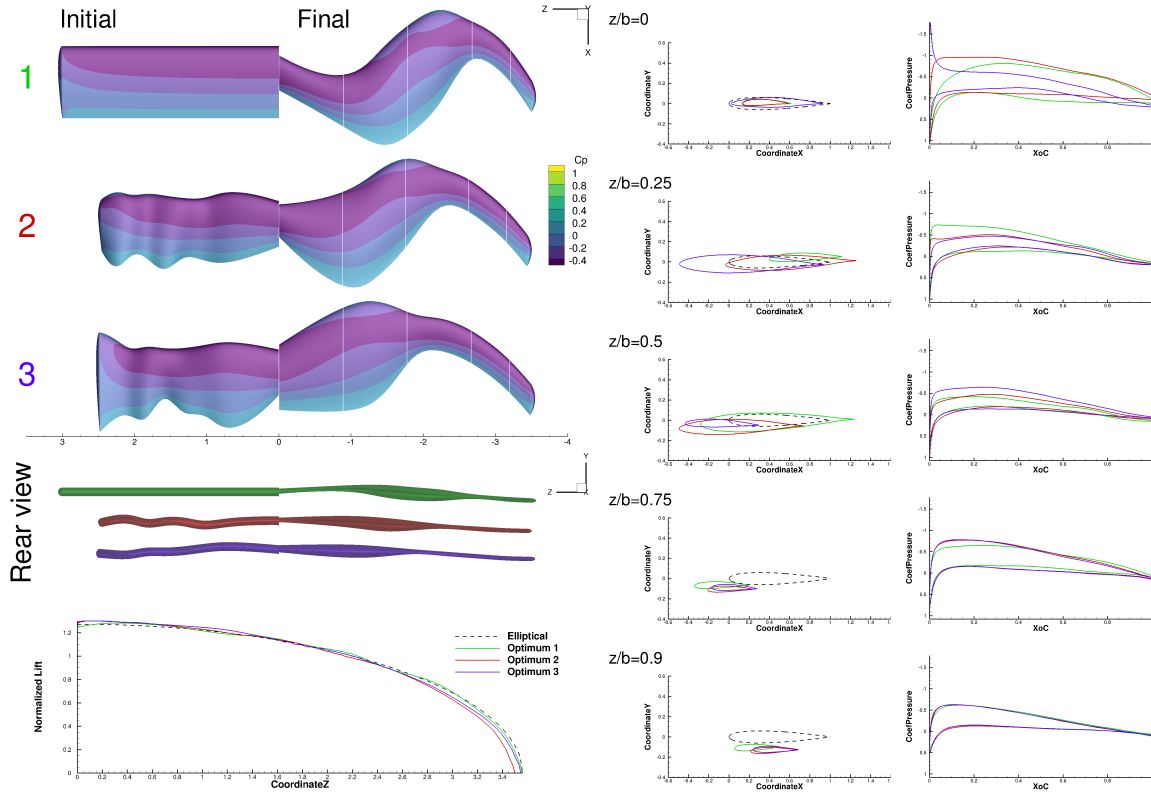


Figure 18: Three local minima found with all variables active.

Table 7: Full optimization problem results. Percentage change is referenced to the baseline wing.

Case	Notes	Optimum	C_L	C_D (cts)	$\% \Delta C_D$	$\% \Delta C_{D,v}$	$\% \Delta C_{D,p}$
1	All variables	1	0.2625	128.358	-10.87	-1.3	-21.7
		2	0.2625	128.847	-10.53	-1.0	-21.3
		3	0.2625	128.888	-10.50	-0.7	-21.6
2	With root bending moment constraint	1	0.2625	131.903	-8.41	-2.8	-14.7
3	With monotonic chord constraint	1	0.2625	128.157	-11.01	-0.8	-22.5
		2	0.2625	128.116	-11.04	-0.3	-23.1
		3	0.2625	128.115	-11.04	-0.3	-23.1
4	With linear chord and sweep distributions	1	0.2625	134.011	-6.95	0.6	-15.4
		2	0.2625	134.011	-6.95	0.6	-15.4
		3	0.2625	134.010	-6.95	0.6	-15.4

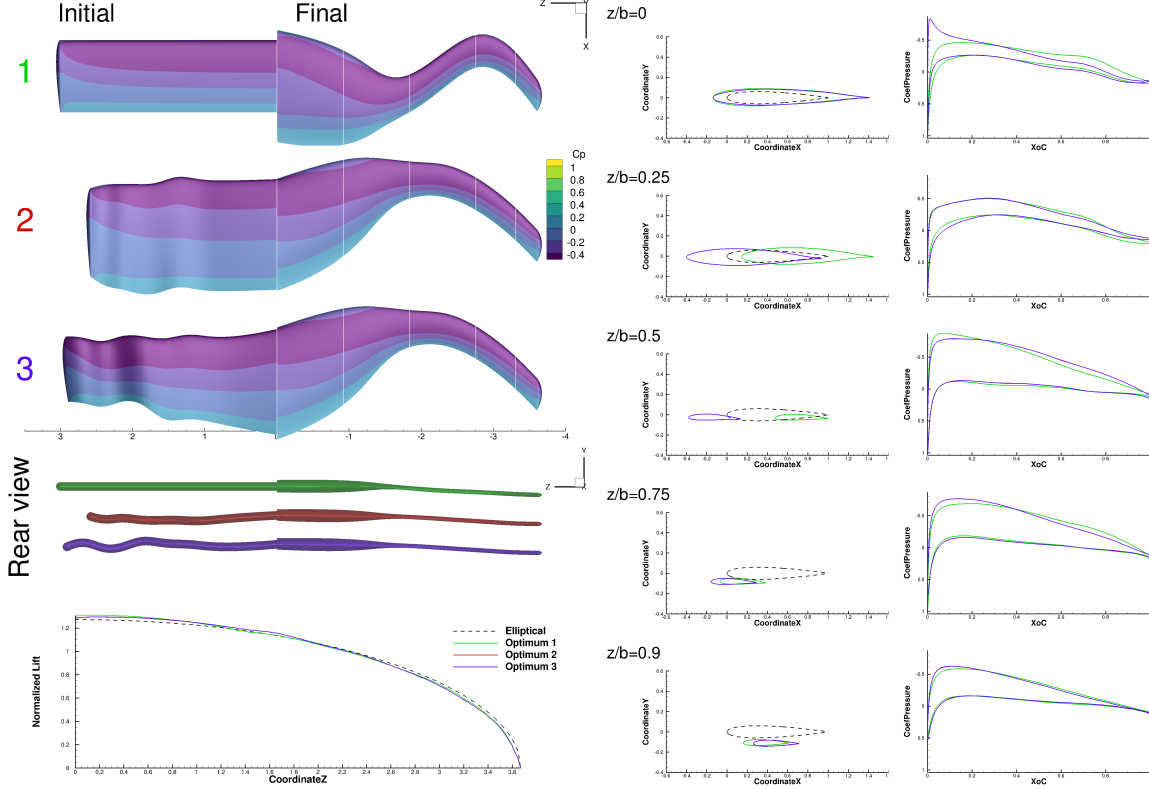


Figure 19: Two local minima found when monotonically decreasing chord is enforced.

better reduction in pressure drag for the Case 3 optima compared with the Case 1 optima. These results corroborate our previous assertion that minimizing skin friction drag with respect to the chord distribution is a multimodal problem. The optima found in Case 3 were present in the design space of Case 1, but the optimizer was unable to find them because of the multimodality due to the chord variables. Importantly, failing to find the optimal chord distribution prevented the optimizer from maximizing span and achieving the optimal aspect ratio. The major difference between the two optima in Case 3 is the sweep distribution. Optimum 1 follows a zig-zag pattern, sweeping back, then forward, then back again at the tip. Optima 2 and 3 have a more gradual sweep distribution, sweeping forward until midspan and then back to the tip. Interestingly, the same sweep patterns are observed in Fig.18. The distribution of anhedral is nearly identical across the optimized designs, but it varies only slightly from a flat wing, indicating the lack of benefit to be gained from a winglet at low C_L .

While these optimized wings certainly perform better than the baseline wing, there is no disputing the fact that they are unconventional—possibly too unconventional to be practical. The final optimization case attempts to address the tradeoff between performance and practicality. For this final case, we add a constraint that the chord and sweep distribution must vary linearly along the span of the wing. In three different runs, the optimizer converges on a single design, shown in Fig.20. This design appears much more conventional, yet it suffers a 4% increase in drag compared with the optimized designs from Cases 1 and 3. The optimum design has maximum sweep and minimum tip chord. The span upper bound is not reached, probably due to the need to maintain a relatively large root chord in order to not cause a greater increase in skin friction drag.

5 Conclusion

Wing aerodynamic optimization with respect to planform variables involves a strong trade-off between induced and parasite drag. The inclusion of a viscous drag model in the optimization is necessary to fully account for this trade-off and obtain physically meaningful results. When minimizing drag with respect to

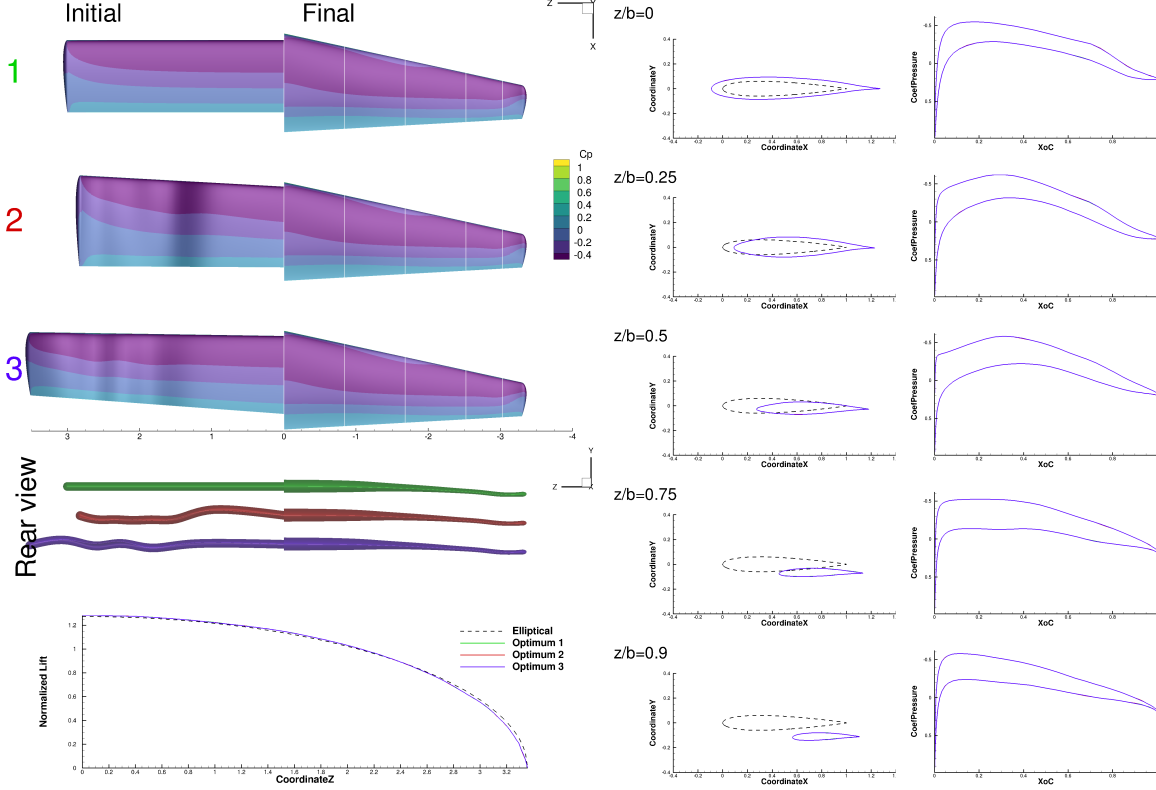


Figure 20: One optimal design found when chord and sweep distributions are constrained to be linear.

chord distribution at low C_L values, this trade-off causes the optimizer to form nonintuitive wavy chord distributions that minimize skin friction drag at the cost of a perfectly elliptical lift distribution. This trade-off introduces multimodality into the design space, which we explained by analyzing a canonical flat plate with zero lift. In this extreme case, the optimal chord distribution is governed by the motivation to increase the chord as much as possible to reduce the skin friction drag per unit area while meeting the wing area constraint, but multiple chord distributions yield the minimum drag solution. As C_L increases, this multimodality is attenuated by the increasing relative importance of induced drag because the induced drag is also coupled with the chord distribution.

The benefit of nonplanarity is strongly correlated with the drag trade-off as well. When optimizing with respect to dihedral variables, the optimizer forms a winglet, for both $C_L = 0.5$ and $C_L = 0.8$. The optimizer also displaces the main wing vertically in the opposite direction of the winglet to maximize the vertical extent of the wing and thus the vorticity sheet. This trend is amplified for $C_L = 0.8$ compared with $C_L = 0.5$. On the other hand, for a low C_L condition, the decrease in induced drag due to nonplanarity does not outweigh the increase in parasite drag due to the additional wetted area. Thus, for high C_L conditions, dihedral variables add multimodality to the design space. We find that a down-turned winglet reduces drag more than an upturned winglet, but both are local minima. A wing optimized with respect to both chord and dihedral variables appears to linearly combine the trends found in the separate cases: at low C_L the planform is wavy with indiscernible dihedral, and at high C_L the wing tapers to a winglet.

For each lift condition there is an optimal aspect ratio. When induced drag is the principal drag component, this optimal aspect ratio may not be possible due to span restrictions. However, we find that at $C_L = 0.2625$, the optimal aspect ratio is attainable within the bounds of the span variable, such that the optimizer does not extend the wing to the maximum possible span. Finally, when the optimizer has freedom to control sweep, it almost always converges to a wing with a sharply swept wingtip.

It is crucial to consider viscous effects when the geometric parametrization allows the development of a nonplanar wing. In such cases, reducing induced drag is an incentive to increase dihedral and form a winglet, which also increases the total surface of the wing. Viscous effects must be considered in order to balance the

tradeoff between induced drag and parasite drag regarding the formation of a winglet. All of these results strongly support the use of viscous analysis over inviscid analysis in realistic wing design optimization.

Although we did find evidence of multimodality for various parametrizations, we do not find these instances to be an impediment to mainstream use of gradient-based aerodynamic shape optimization of wings. The use of proper physics modeling and the application of practical design constraints enabled us to arrive at a single optimal design for the problem considered in this paper. The insights gleaned from these results will help other researchers gain a better understanding of the wing design problem and formulate their problems to avoid possible sources of multimodality. Our experience is that in practical design problems, the number of local minima is automatically reduced merely by the standard constraints and requirements of the design problem. As stated in the introduction, we did not attempt to, nor did we, conduct an exhaustive search for all of the local minima in the design space. Our study confirms our opinion that it is not necessary or economical to engage in such activity for each new design problem.

There are a few avenues for future work to build on these results. Since the optimal wing design is strongly dependent on C_L , it would be interesting to study the effect of optimizing with respect to multiple flight conditions. This would give a more realistic representation of real-world wing performance and should produce a more versatile wing. Additionally, a structural model should be considered, since real-world wing design requires aerostructural analysis and design trade-offs.

6 Acknowledgements

This research was partially funded by Embraer S.A. The authors would like to thank Gaetan Kenway and John Jasa for their development of the MACH and OpenAeroStruct frameworks, respectively.

References

- [1] Anderson, J. D., *Introduction to Flight*, McGraw-Hill, 2011.
- [2] Martins, J. R. R. A., and Lambe, A. B., “Multidisciplinary Design Optimization: A Survey of Architectures,” *AIAA Journal*, Vol. 51, No. 9, 2013, pp. 2049–2075. doi:[10.2514/1.J051895](https://doi.org/10.2514/1.J051895).
- [3] Chittick, I. R., and Martins, J. R. R. A., “An Asymmetric Suboptimization Approach to Aerostructural Optimization,” *Optimization and Engineering*, Vol. 10, No. 1, 2009, pp. 133–152. doi:[10.1007/s11081-008-9046-2](https://doi.org/10.1007/s11081-008-9046-2).
- [4] Hicks, R. M., and Henne, P. A., “Wing Design by Numerical Optimization,” *Journal of Aircraft*, Vol. 15, 1978, pp. 407–412.
- [5] Jameson, A., Martinelli, L., and Pierce, N. A., “Optimum Aerodynamic Design Using the Navier–Stokes Equations,” *Theoretical and Computational Fluid Dynamics*, Vol. 10, No. 1–4, 1998, pp. 213–237. doi:[10.1007/s001620050060](https://doi.org/10.1007/s001620050060).
- [6] Ishimitsu, K., “Aerodynamic design and analysis of winglets,” *Aircraft Systems and Technology Meeting*, 1976, p. 940.
- [7] Pfeiffer, N., “Numerical winglet optimization,” *42nd AIAA Aerospace Sciences Meeting and Exhibit*, 2004, p. 213.
- [8] Jansen, P., Perez, R. E., and Martins, J. R. R. A., “Aerostructural Optimization of Nonplanar Lifting Surfaces,” *Journal of Aircraft*, Vol. 47, No. 5, 2010, pp. 1491–1503. doi:[10.2514/1.44727](https://doi.org/10.2514/1.44727).
- [9] Martins, J. R. R. A., Alonso, J. J., and Reuther, J. J., “High-Fidelity Aerostructural Design Optimization of a Supersonic Business Jet,” *Journal of Aircraft*, Vol. 41, No. 3, 2004, pp. 523–530. doi:[10.2514/1.11478](https://doi.org/10.2514/1.11478).
- [10] Lyu, Z., Kenway, G. K. W., and Martins, J. R. R. A., “Aerodynamic Shape Optimization Investigations of the Common Research Model Wing Benchmark,” *AIAA Journal*, Vol. 53, No. 4, 2015, pp. 968–985. doi:[10.2514/1.J053318](https://doi.org/10.2514/1.J053318).

- [11] Kenway, G. K. W., and Martins, J. R. R. A., “Multipoint High-Fidelity Aerostructural Optimization of a Transport Aircraft Configuration,” *Journal of Aircraft*, Vol. 51, No. 1, 2014, pp. 144–160. doi:[10.2514/1.C032150](https://doi.org/10.2514/1.C032150).
- [12] Brooks, T. R., Kennedy, G. J., and Martins, J. R. R. A., “High-fidelity Multipoint Aerostructural Optimization of a High Aspect Ratio Tow-steered Composite Wing,” *Proceedings of the 58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, AIAA SciTech Forum*, Grapevine, TX, 2017. doi:[10.2514/6.2017-1350](https://doi.org/10.2514/6.2017-1350).
- [13] Burdette, D. A., and Martins, J. R. R. A., “Impact of Morphing Trailing Edge on Mission Performance for the Common Research Model,” *Journal of Aircraft*, 2018. doi:[10.2514/1.C034967](https://doi.org/10.2514/1.C034967), (In press).
- [14] Lyu, Z., and Martins, J. R. R. A., “Aerodynamic Design Optimization Studies of a Blended-Wing-Body Aircraft,” *Journal of Aircraft*, Vol. 51, No. 5, 2014, pp. 1604–1617. doi:[10.2514/1.C032491](https://doi.org/10.2514/1.C032491).
- [15] Mader, C. A., Kenway, G. K., Martins, J. R. R. A., and Uranga, A., “Aerostructural Optimization of the D8 Wing with Varying Cruise Mach Numbers,” *18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, American Institute of Aeronautics and Astronautics, 2017. doi:[10.2514/6.2017-4436](https://doi.org/10.2514/6.2017-4436).
- [16] Gagnon, H., and Zingg, D. W., “Two-Level Free-Form Deformation for High-Fidelity Aerodynamic Shape Optimization,” *12th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference and 14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 2012.
- [17] Hicken, J. E., and Zingg, D. W., “Aerodynamic Optimization Algorithm with Integrated Geometry Parameterization and Mesh Movement,” *AIAA Journal*, Vol. 48, No. 2, 2009, pp. 400–413.
- [18] Kenway, G. K., Kennedy, G. J., and Martins, J. R. R. A., “A CAD-Free Approach to High-Fidelity Aerostructural Optimization,” *Proceedings of the 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, Fort Worth, TX, 2010. doi:[10.2514/6.2010-9231](https://doi.org/10.2514/6.2010-9231).
- [19] Chernukhin, O., and Zingg, D. W., “Multimodality and Global Optimization in Aerodynamic Design,” *AIAA Journal*, Vol. 51, No. 6, 2013, pp. 1342–1354. doi:[10.2514/1.j051835](https://doi.org/10.2514/1.j051835).
- [20] Streuber, G. M., and Zingg, D. W., “Investigation of multimodality in aerodynamic shape optimization based on the Reynolds-Averaged Navier-Stokes equations,” *Proceedings of the 18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Denver, CO, 2017.
- [21] Poole, D. J., Allen, C. B., and Rendall, T., “Global Optimization of Multimodal Aerodynamic Optimization Benchmark Case,” *35th AIAA Applied Aerodynamics Conference*, 2017, p. 4365.
- [22] Lyu, Z., Kenway, G. K., Paige, C., and Martins, J. R. R. A., “Automatic Differentiation Adjoint of the Reynolds-Averaged Navier-Stokes Equations with a Turbulence Model,” *21st AIAA Computational Fluid Dynamics Conference*, San Diego, CA, 2013. doi:[10.2514/6.2013-2581](https://doi.org/10.2514/6.2013-2581).
- [23] Kenway, G. K. W., Secco, N., Martins, J. R. R. A., Mishra, A., and Duraisamy, K., “An Efficient Parallel Overset Method for Aerodynamic Shape Optimization,” *Proceedings of the 58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, AIAA SciTech Forum*, Grapevine, TX, 2017. doi:[10.2514/6.2017-0357](https://doi.org/10.2514/6.2017-0357).
- [24] Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A., “Scalable Parallel Approach for High-Fidelity Steady-State Aeroelastic Analysis and Derivative Computations,” *AIAA Journal*, Vol. 52, No. 5, 2014, pp. 935–951. doi:[10.2514/1.J052255](https://doi.org/10.2514/1.J052255).
- [25] Jasa, J. P., Hwang, J. T., and Martins, J. R. R. A., “Open-source coupled aerostructural optimization using Python,” *Structural and Multidisciplinary Optimization*, Vol. 57, No. 4, 2018, pp. 1815–1827. doi:[10.1007/s00158-018-1912-8](https://doi.org/10.1007/s00158-018-1912-8).

- [26] Gray, J. S., Hwang, J. T., Martins, J. R. R. A., Moore, K. T., and Naylor, B. A., “OpenMDAO: An open-source framework for multidisciplinary design, analysis, and optimization,” *Structural and Multidisciplinary Optimization*, 2019. (Accepted subject to minor revisions).
- [27] Gill, P. E., Murray, W., and Saunders, M. A., “An SQP algorithm for large-scale constrained optimization,” *Society for Industrial and Applied Mathematics*, Vol. 47, No. 1, 2005. URL <http://www.stanford.edu/group/SOL/papers/SNOPT-SIGEST.pdf>.
- [28] Burdette, D. A., Kenway, G. K. W., and Martins, J. R. R. A., “Aerostructural design optimization of a continuous morphing trailing edge aircraft for improved mission performance,” *17th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Washington, D.C., 2016. doi:10.2514/6.2016-3209.
- [29] Lee, C., Koo, D., Telidetzki, K., Buckley, H., Gagnon, H., and Zingg, D. W., “Aerodynamic Shape Optimization of Benchmark Problems Using Jetstream,” *Proceedings of the 53rd AIAA Aerospace Sciences Meeting*, American Institute of Aeronautics and Astronautics (AIAA), 2015. doi:10.2514/6.2015-0262.
- [30] Bisson, F., and Nadarajah, “Adjoint-based aerodynamic optimization of benchmark problems,” *53rd Aerospace Sciences Meeting*, 2015.
- [31] Poole, D. J., Allen, C. B., and Rendall, T., “Control point-based aerodynamic shape optimization applied to aiaa adodg test cases,” *53rd AIAA Aerospace Sciences Meeting*, 2015, p. 1947.
- [32] Munk, M. M., “The Minimum Induced Drag of Aerofoils,” Tech. Rep. NACA-TR-121, 1923.
- [33] Smith, S. C., and Kroo, I. M., “Computation of induced drag for elliptical and crescent-shaped wings,” *Journal of aircraft*, Vol. 30, No. 4, 1993, pp. 446–452.
- [34] Blasius, H., “Grenzschichten in Flüssigkeiten mit kleiner Reibung,” *Zeitschrift für angewandte Mathematik und Physik*, Vol. 56, 1908, pp. 1–37.
- [35] Khosravi, S., and Zingg, D. W., “Aerostructural Perspective on Winglets,” *Journal of Aircraft*, 2017.