A Parametric Study of Multimodality in Aerodynamic Shape Optimization of Wings

Gregg M. Streuber*
David W. Zingg†
University of Toronto, Toronto, Ontario, Canada

A parametric study of multimodality in wing optimization is undertaken utilizing a gradient-based multistart method. The lift-constrained drag minimization of a wing is performed using between 17 and 33 initial geometries in subsonic and transonic viscous flows with varying degrees of freedom and under a variety of constraints. In nearly every examined case, including the ADODG CRM case, multimodality of some degree is found. The cross-sectional optimization of a wing is found to be unimodal to somewhat multimodal in both subsonic and transonic flows, depending on which degrees of freedom are permitted. Permitting large-scale planform deformations and non-planar geometries is shown to produce clear and significant multimodality in all examined cases. Requiring straight leading and trailing edges can significantly reduce, but not fully eliminate, multimodality. Other examined constraints, such as linear taper, linear twist, minimum thickness, and minimum pitching moment are shown to have lesser effects on multimodality. Multimodality is ultimately found to be an inherent characteristic of most wing optimization design spaces, and it is concluded that gradient-based optimization based on a single initial geometry may not be sufficient to ensure global optimality even in tightly constrained practical problems.

I. Introduction

The spectre of anthropogenic climate change, unstable fuel prices, and diminishing returns from the continued optimization of conventional aircraft designs are driving a search for dramatically more efficient new aerospace technologies.[1] A major pillar of this effort is research into higher-efficiency unconventional aircraft configurations, an undertaking that is increasingly reliant on high-fidelity aerodynamic shape optimization.

When applying numerical optimization to any problem an understanding of the multimodality, if any, of the design space is critical. This is because the presence or extent of multimodality plays a significant role in determining the proper tools and approaches for solving a given problem. Gradient-based optimization algorithms utilizing the adjoint method[2] are relatively cost-effective, but are stymied by even low degrees of multimodality, and any confidence that the obtained solution is the global optimum only extends so far as the confidence that the design space is unimodal. On the other hand, gradient-free approaches like genetic algorithms[3] are quite robust with regard to even highly multimodal problems, but require a greater investment of computational resources than comparable gradient-based methods,[4] even if one considers potential cost savings from the use of hierarchical genetic algorithms,[5] surrogate methods,[6–8] or hybrid algorithms.[9,10] It has been demonstrated that multimodality can exist in aerodynamic shape optimization problems,[11] and the importance of understanding multimodality in this context is underscored by the recent inclusion of a multimodal problem in the Aerodynamic Design and Optimization Discussion Group (ADODG) test suite at the AIAA Aviation 2017 conference.[12–15]

The objective of this paper is to present a detailed study of the relationship between various problem parameters and multimodality in aerodynamic shape optimization of wings. This includes variations in available degrees of freedom, constraints, as well as Mach number, with particular attention paid to differences

*Ph.D. Candidate, University of Toronto Institute for Aerospace Studies, Toronto Ontario (gregg.streuber@mail.utoronto.ca)
†University of Toronto Distinguished Professor of Computational Aerodynamics and Sustainable Aviation, Director, Center for Sustainable Aviation, Director, Centre for Computational Science and Engineering, and Associate Fellow AIAA (dwz@oddjob.utias.utoronto.ca)

in outcome between subsonic and transonic flight regimes. This study is performed using the NASA common research model (CRM) wing-only geometry as the baseline. The viscous optimization of this geometry has been studied by numerous authors, including - but certainly not limited to - Koo and Zingg,[16] Yu et al.,[17] Lee et al.,[18] Shi-Dong et al.,[19] and Lyu et al.[20] Several of these studies[16,17,20] directly address the question of multimodality, generally in the context of the transonic, viscous optimization of twist and taper, the consensus being that multimodality under these conditions is minimal to nonexistent. On the other end of the spectrum, the existence of multimodality in the aerodynamic shape optimization of a wing in general is demonstrated in several previous works.[11,15,21] Between these extremes must lie a region in which multimodality changes from significant to insignificant and it is a goal of this work to determine the boundaries of these regions within the design space. This is accomplished by utilizing the free-form deformation[22] (FFD) based gradient-based multistart (GBMS) algorithm the authors published previously[15] to thoroughly explore the design space, specifically addressing the causes and extent of multimodality in the optimization of an aircraft wing under varying degrees of freedom, constraints, and operating conditions. This is of critical importance in determining under what conditions local optimizers are sufficient to explore a problem, and when more expensive global optimization schemes must be leveraged to provide reliable and thorough results.

In any practical aerodynamic shape optimization problem only a finite degree of convergence is possible. This allows the possibility of apparent local optima that would ultimately have converged to the same optimum if further convergence could be achieved. Therefore some judgement is required to distinguish these apparent local optima from actual local optima.

II. Methodology

A. Gradient-Based Multistart Algorithm

Gradient-based multistart is a method wherein a sampling algorithm is used to deform a single baseline geometry into a large sample of the design space, with each sample point being a unique geometry that is then optimized in parallel using a gradient-based method. The samples themselves are generated using a Sobol sampling method, the specific implementation being based on Algorithm 659,[23] which uses more primitive polynomials than other algorithms, and a Gray code implementation proposed by Antonov and Saleev.[24] This yields a method that is more robust in the face of multimodality than a single gradient-based optimization but less expensive than gradient-free methods, particularly for "moderately" multimodal problems with $\mathcal{O}(10^1)$ or fewer local optima.[11] This balance between cost and robustness makes GBMS ideal for exploring new design spaces, particularly when the objective is to ascertain the extent of multimodality in the system in order to determine the appropriate optimization algorithm for the particular problem class at hand.

While the GBMS label was first coined by Chernukhin and Zingg, variations on this idea are present in the literature under a number of different names.[25–29] Chernukhin and Zingg developed specialized linear constraints which ensure the generated samples can be easily tailored to produce feasible geometries within the geometric bounds desired by the user. In our previous work[15] we expanded on this method, utilizing FFD geometry control to generalize these constraints to a geometry-independent form, producing an algorithm that can be applied to any geometry or class of problem.

B. Free-Form Deformation

A common analogy is to think of FFD[22] as embedding an object within a rubbery block. If one deforms the rubber block, the underlying geometry is deformed as well. Computationally, this is accomplished by embedding an underlying parameterization, here B-spline surfaces, inside another B-spline volume, referred to as the "FFD volume". Each point of the underlying parameterization is assigned a parameteric position within the FFD volume, and as the volume is deformed, the positions of the underlying control points are re-evaluated to keep their parametric positions constant.

The two-level FFD implementation used in this work was developed by Gagnon and Zingg,[30] and is notable for its use of an underlying B-spline parameterization as the embedded points, rather than the computational grid nodes themselves. This control scheme also includes axial control, which adds a NURBS curve – referred to as the "axial" curve – which is used to drive large scale deformations. In this approach, the FFD volume is divided into planar slices referred to as cross sections, constrained to be locally perpendicular to the axial curve and offering sectional control through various transformations of the cross section and the manipulation of two rows of control points on the upper and lower surfaces of the FFD volume. The axial curve is controlled through a series of axial control points; deforming the curve in this way produces large scale deformations of the entire volume, while leaving the local coordinate systems untouched. There are a total of six types of geometric degrees of freedom available in this system: twist, taper and section (cross-sectional) and sweep, span and dihedral (axial).

C. Aerodynamic Shape Optimization Framework

Optimization is accomplished using Jetstream,[30–32] which couples a three-dimensional, finite-difference, structured, multiblock, parallel, implicit flow solver ("Diablo") that can be applied to either the Euler[33] or Reynolds-averaged Navier Stokes (RANS)[34] equations, with a gradient-based optimization code built around the SQP Sparse Nonlinear OPTimizer (SNOPT),[35] and an integrated linear-elasticity mesh movement algorithm.[31]

In Diablo, the equations are discretized via second-order summation-by-parts operators with scalar or matrix numerical dissipation; boundary conditions and block interfaces are enforced with simultaneous approximation terms. The resulting equations are solved using a parallel Newton-Krylov-Schur algorithm, consisting of an approximate-Newton phase which obtains the initial iterate for a subsequent inexact-Newton phase. Both stages solve the large linear systems with GMRES, a Krylov iterative solver with approximate-Schur preconditioning. The RANS equations are closed with the Spalart-Allmaras one-equation turbulence model.[36]

Gradients provided to the optimizer are calculated using the adjoint method, so the cost is nearly indepedent of the number of design variables. Optimality is ensured by enforcing the Karush-Kuhn-Tucker optimality conditions.[37] The preconditioned conjugate gradient method is applied to the mesh adjoint system, while a flexible variant[38] of the GCROT Krylov method[39] solves the flow adjoint system.

D. Test Structure

The objective of this work is to offer a detailed exploration of the circumstances under which multimodality becomes a concern and to what extent during the optimization of a wing where the flow is governed by the RANS equations. To accomplish this, a number of tests are run using identical baseline geometries and numerical meshes, specifically the NASA CRM wing-only case, described in more detail in Section III.A. In each case, a parameter is varied – generally the number and type of degrees of freedom, the linear or nonlinear constraints, or the operating conditions – and a large sample is generated using the algorithm from Streuber and Zingg.[15] Seventeen tests, 16 samples and the baseline geometry, are then optimized in parallel to either failure or until sufficient convergence has been achieved; the resulting geometries are examined to determine how many local optima have appeared. If multiple local optima are discovered, they are noted and the test is stopped. Otherwise, a further 16 cases are run, and if once again no local optima appear, the problem is considered to be unimodal for the purposes of this study. We are not seeking here to locate the global optimum; doing so with confidence for an unknown design space would require a substantially larger number of initial geometries. A sample of this size is sufficient to show with some confidence whether or not a given problem is multimodal and to permit a discussion of the relative degrees of multimodality between separate problems.

Two indicators are used to determine whether two geometries represent distinct optima: performance and geometric variation. For the first, a material difference in performance must be apparent between the geometries, generally defined as something on the order of one or more drag counts. For the second indicator, the two geometries should be substantially different; this is done by comparing the root mean square difference between the surface control points of each geometry, with differences greater than approximately 5×10^{-2} being considered "substantial", though judgement is exercised in this manner by visually inspecting the geometries as well. If either of these criteria are met, the geometries in question are considered to be two discrete optima. It is not enough to solely check the geometric variation, as there are numerous scenarios in which small geometric differences can produce material variations in performance – changes to the spanwise twist distribution being a notable example. Once all optima have been located, several metrics, discussed below, can be used to draw conclusions on the degree and importance of multimodality within the design space, and the impact of the tested parameters on those values.

E. Multimodality Metrics

The overall amount of multimodality in a design space can be fully described through several quantities. The first is the number of local optima, this is the number of distinct local optima that the initial geometries ultimately converged to. The second is the performance range, defined as the difference in the performance between the best and worst local optima. The third is the dominance, the percentage of converged geometries which arrived at the most frequent – though not necessarily best performing – local optimum. These provide a clear high-level picture of the degree of multimodality in a design space and the ranges within which the local optima, if any, lie. However, they give us little to no information on how the local optima are distributed within the design space, which is critical to determining the best approach to the solution of a given problem.

Consider two design spaces, each with 10 local optima, a minimum drag of 180 drag counts, and a performance range of 15 drag counts. In design space A, nine of the local optima are clustered near the global optimum, with only a single non-dominant optimum at the upper performance range. In design space B, the optima are more scattered throughout the design space, with one or two near the minimum and the rest evenly distributed in terms of performance. If one is focused on locating the minimum drag in the design space, whether you are in design space A or B will significantly change the ideal approach to the problem: in A the vast majority of initial positions will converge to near the optimal design, and so a handful of initial geometries should be sufficient to have confidence that the global optimum has been reached. In B the likelihood of finding a significantly inferior optimum is much higher, and so a reliable optimization would require a commensurately larger sample size or – depending on the problem – considering a move to hybrid or gradient-free algorithms. Alternatively, suppose we are focused on locating all unique geometries within these design spaces – as perhaps these novel geometries offer other, non-aerodynamic benefits that may outweigh their inferior drag values. In this scenario, A is in fact the more difficult design space, as an insufficiently large sample would make it easy to miss the outlying optima, only capturing the likely highly similar geometries near the global optimum. Conversely, the more evenly distributed nature of B makes it more likely that a comparatively small sample will capture all the major geometric trends in the design space.

These conflicting priorities – sometimes wishing to bypass all outliers, and other times needing to locate as many as possible – are reflected in two supplementary multimodality metrics which seek to quantify the distribution of optima within the design space. The first is the performance metric, P_m, defined as

$$P_{\mathrm{m}} = 100 \frac{\sum n_i |\mathcal{J}_i - \mathcal{J}_{\mathrm{best}}|}{\mathcal{J}_{\mathrm{best}} \sum n_i}, \tag{1}$$

where $\mathcal{J}_i$ is the objective function at the i^{th} local optimum, $\mathcal{J}_{\mathrm{best}}$ is the best performance achieved at any optimum, and n_i is the number of initial geometries which converged to the i^{th} local optimum. The equation is scaled by a factor of 100 for convenience and readability, so that the resulting values will be on the order of $\mathcal{O}(10^{-1})$ to $\mathcal{O}(10^{0})$. This quantifies not just the number of local optima, but how common each is and how their performance differs from the best optimum. Generally, a larger value of P_m indicates a design space in which the best performing geometry is increasingly difficult to locate and suggests that confidence in the obtained solution will require a GBMS method with a progressively larger sample size, or moving to a hybrid or gradient-free approach.

The second metric is the geometric metric, denoted as G_m and defined as

$$\begin{aligned} G_{\mathrm{m}} &= 100 \frac{1}{\sum \frac{1}{n_i}} \sum \frac{R_i}{n_i}, \\ R_i &= \sqrt{\sum (x_{i,j} - x_{\mathrm{dom},j})^2}, \end{aligned} \tag{2}$$

as before n_i is the number of initial geometries which converged to the i^{th} local optimum, and R_i is the root mean square difference between the i^{th} and the most dominant local optimum, or the optimized baseline if no optimum is most dominant. The values of x define the three-dimensional coordinates of each point on the surface of the geometry. An important detail is that G_m and P_m are inversely weighted with respect to n_i. Therefore, P_m penalizes dominant, inferior local optima which are likely to impede progress towards the global optimum, while G_m penalizes non-dominant, geometrically differentiated local optima which are likely to be difficult to locate. A large G_m value suggests that there are many unique geometries in difficult-to-access regions of the design space and that therefore locating all major geometric modes will require a very thorough exploration of the design space.

III. Results

All tests are performed using the same baseline geometry, presented below. For clarity, within this paper "baseline" is taken to refer to the fundamental initial geometry and problem definition that the study is built around. This problem definition is varied throughout this paper to study various parameters, and the baseline geometry itself is modified by the sampling algorithm to produce the initial geometries which are then optimized. The first study looks for multimodality within the fundamental CRM ADODG test case, this is followed by an investigation of the impact of degrees of freedom on multimodality, then the effect of Mach number is studied, and finally an exploration of how constraints affect multimodality is undertaken.

A. Problem Definition

The problem and baseline geometry here are derived from the CRM wing-body configuration from the Fifth Drag Prediction Workshop,[40] with the fuselage removed and the wing scaled by the mean aerodynamic chord. The problem definition itself naturally varies, as it is the purpose of this paper to examine the impact on multimodality of variations in the problem definition. However, the fundamental core of the problem, defined below, remains fixed. Specifically, this is the RANS-based optimization of an aircraft wing. The Mach and Reynolds numbers vary from case to case and will be clearly stated where relevant. Volume is constrained to be no less than the initial volume of 0.2617 cubed reference units, and projected area must equal the initial value of 3.407 square units, while the lift coefficient is constrained to $C_L = 0.5$. Formally, the optimization problem can be written as:

$$\begin{aligned} \text{minimize} \quad & C_D \\ \text{w.r.t.} \quad & v \\ \text{subject to} \quad & C_L = 0.5 \\ & S = 3.407 \\ & V \geq 0.2617 \end{aligned} \tag{3}$$

It should be noted that this is a very lightly constrained problem, and intentionally so. Generous upper and lower bounds are set for each design variable to maintain a bounded design space; unless otherwise stated there are no other significant constraints on the design variables during optimization.

B. Geometry Control

The baseline geometry is depicted in Figure 1a. Control is provided by two FFD volumes, which meet at the crank. The inboard FFD volume is controlled by three cross sections, while the larger outboard volume is divided into five. Large scale deformations of both volumes are driven by one of two axial curves, each controlled by five axial control points. This control scheme remains fixed throughout all tests, though the available degrees of freedom do vary.

C. Optimization Grid

A single moderately dense grid is used for all analyses, as the focus is on the number of local optima obtained from each test and as such high numerical accuracy is not needed in the values of the force and moment coefficients. Relevant parameters for this grid are provided in Table 1.

D. CRM ADODG Case

Prior to examining more open-ended problems, we first study the ADODG CRM test case. This is performed at a Mach number of 0.85 and a Reynolds number of 5 million, using only twist and sectional control design variables. In addition to the constraints listed in the fundamental problem definition, a pitching moment constraint is enforced, requiring that $C_M \geq -0.17$. In accordance with the previously outlined test procedure,

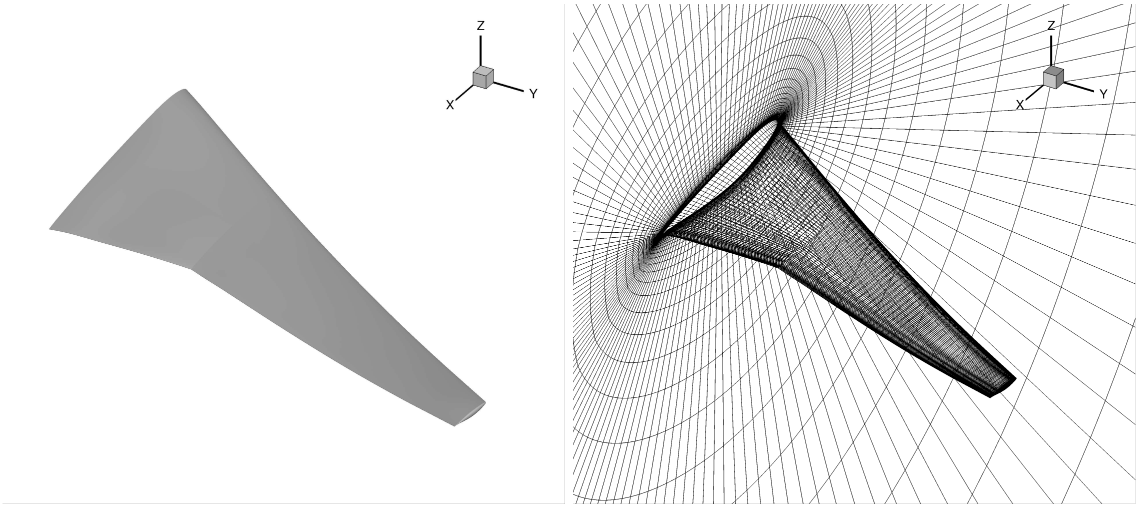

(a) Baseline geometry (b) Baseline mesh

Figure 1: Geometry and mesh for baseline geometry

Table 1: Common Research Model mesh parameters

Grid	Nodes	Off-wall Spacing (mean aerodynamic chord)	Transonic y^+	Subsonic y^+
L0	925,888	2.19×10^{-6}	0.34	0.27

a total of 17 initial geometries were optimized, 16 of which converged successfully. The results of this study are provided in Table 2. We find this problem to be clearly multimodal, though highly dominated by the apparent global optimum. Four of the five local optima are within 2 drag counts of each other, with a single outlying optimum accounting for the 9.6 drag count performance range.

The twist distribution for each local optimum is plotted in Figure 2a and shows the presence of at least three distinct twist distributions between the five identified local optima. The orange curve corresponds to the worst performing local optimum. The spanwise lift distributions are also plotted, in Figure 2b; all five optima produce similar and nearly-elliptical lift distributions. This indicates that the combination of twist and section control permits the optimizer multiple avenues to produce an optimal lift distribution, with varying results vis-a-vis non-induced drag. Regarding the lift distributions the orange curve is once again the clearest outlier; though the fact that this optimum is the furthest from an ideal elliptical lift distribution is not surprising given its inferior performance.

Earlier CRM studies[32] suggested the possibility of multimodality in this problem, though more recent work[16,17,20] has concluded that the design space is unimodal. While we conclude that the problem permits multiple local optima, the design space appears heavily dominated by a single particular optimum. Twelve of the 16 successfully converged geometries, including the baseline, converged to the best optimum. The remaining four geometries each converged to a distinct local optimum. Any study which does not explore the design space in sufficient breadth would be quite likely to miss these local optima. While the study of the basic CRM wing-only case yielded a multimodal design space, this represents just one of a plethora of wing optimization cases commonly examined, either in industrial or research settings. These include varying degrees of freedom, flow conditions, and constraints; one must be careful extrapolating the results from this single case to wing optimization in general. To better understand multimodality in aerodynamic shape optimization, we now examine how multimodality is impacted by these various parameters.

E. Transonic Degree of Freedom Study

To explore the impact of variable degrees of freedom on multimodality we perform a study with the CRM wing as a baseline while permitting the optimizer to utilize different combinations of degrees of freedom at a

Table 2: ADODG CRM Multimodality Study Results

Successful Tests	Number of Optima	Best (drag counts)	Range (drag counts)	Dominance	P_m	G_m
16	5	196	9.6	75%	0.44	0.34

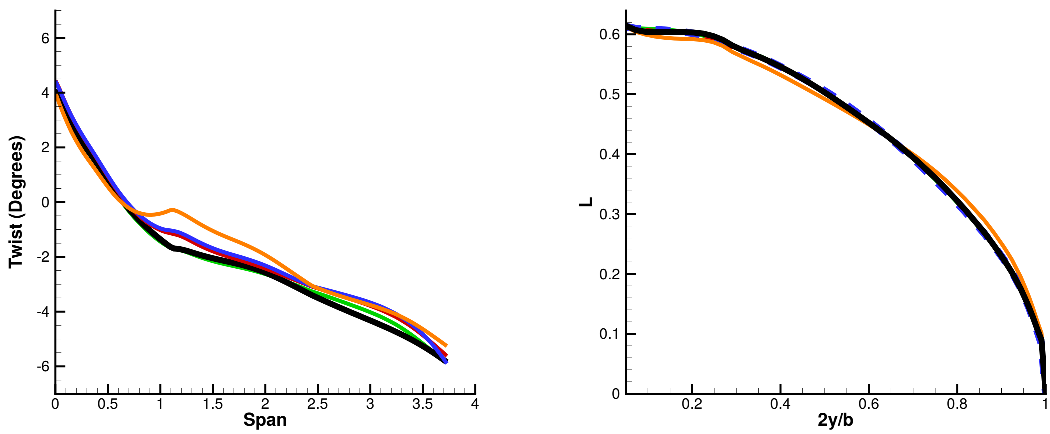

(a) Twist distributions

(b) Lift distributions

Figure 2: Spanwise lift and twist distributions for ADODG CRM case. Best optimum in black.

constant Mach number of 0.85 and Reynolds number of 5 million. Beginning with relatively small, practical problems we gradually add freedom, moving towards increasingly exploratory problems. Only the standard lift, projected area, and minimum volume constraints are enforced. Neither the pitching moment constraint, nor any other significant geometric constraints are enforced on these problems. As such, this is a very loosely constrained problem and the multimodality found here can best be interpreted as an upper bound on the multimodality present for a given combination of design variables, with any additional constraints or requirements likely leading to a commensurately lower degree of multimodality. The impact of various constraints is studied in Section G.

Seven sub-problems are solved by varying the type and number of degrees of freedom available to the optimizer. Each test here is denoted by a five digit code of the form xxx_xx. The first three digits refer to the cross sectional degrees of freedom (from left to right, twist, taper, section shape), while the latter three refer to the sweep and dihedral axial degrees of freedom. Each digit takes a value of either 1 or 0, indicating whether that degree of freedom is active for the case in question. For example, 001_00 denotes a case where only section shape design variables are active, while 111_01 denotes a case where all cross sectional degrees of freedom are permitted, but dihedral is the only active axial degree of freedom.

The results from all seven cases are tabulated in Table 3, which displays the number of successful tests (the number of attempted tests less any which failed to converge adequately), as well as the number of local optima ultimately located, the best performance achieved in drag counts, and the total performance range from best to worst performing local optima, as well as P_m and G_m values.

Several trends are apparent in this data. As expected, multimodality tends to increase with the number of degrees of freedom. This can be seen through growth in the number of local optima and performance ranges, decreasing values of dominance, and increasing values of P_m and G_m. Across all cases, all local optima are shock-free. Examining the cross-section only cases - 001_00, 011_00, 101_00, 110_00, and 111_00 results (corresponding to different combinations of twist, taper, and section control), shows that under these lightly constrained conditions, even a relatively small amount of freedom provides the potential for multimodality, with three cases – 011_00, 101_00, and 111_00 – producing multiple local minima. In all three of these cases a material difference in performance between local optima is seen; however, the consistently low P_m and

Table 3: Degree of Freedom Multimodality Study Numerical Results. M=0.85, Re=5 million. Each test is identified using a five digit code of the format xxx_xx. From left to right each digit indicates whether twist, taper, section shape, sweep, or dihedral freedom is permitted.

Test Code	Successful Tests	Optima	Best (drag counts)	Range (drag counts)	Dominance	P_m	G_m
001_00	31	1	193	0.0	100%	0.00	0.00
011_00	14	2	188	1.3	86%	0.01	0.75
101_00	17	4	193	2.0	65%	0.16	0.26
110_00	33	1	208	0.0	100%	0.00	0.00
111_00	15	3	187	2.2	73%	0.23	3.92
111_01	12	11	181	7.6	17%	1.35	29.16
111_11	10	9	178	34.8	20%	6.10	74.41

G_m values, particularly for the 011_00 and 101_00 cases, indicates that this multimodality is nevertheless not of a very large degree and that a GBMS algorithm with a small number of initial geometries should be sufficient to fully explore the design space. The 101_00 case differs from the ADODG CRM case only in the absence of a pitching moment constraint, with the 101_00 showing a performance range roughly 25% as large and producing much lower values of P_m and G_m compared to the ADODG case. The results imply that the design space is in fact more difficult to navigate with the pitching moment constraint than without. These differences are entirely due to the single significantly outlying local optimum identified in the ADODG case. Removing this optimum from the results yields a performance range of 2.1 drag counts and P_m and G_m values, respectively, of 0.14 and 0.26, nearly identical to those listed in Table 3 for 101_00.

This presents two likely explanations: first, that this additional optimum is not present in the 101_00 design space, or secondly that it is present in both design spaces but was not located in the 101_00 study. While the constraint study undertaken in Section III.G shows that the predominant impact of increasing constraints is to reduce multimodality, constraints also potentially introduce new trade-offs into the design problem which could in some cases lead to the introduction of new local optima. This is hinted at by the inconsistent relationship between some of the studied constraints in Section III.G and the resulting degree of multimodality. With regard to the second possible explanation, the sample sizes used in this work are not sufficient to ensure an exhaustive study of the design space. For this reason, rare or difficult-to-find local optima may be missed, this is a possibility that will be raised at several points throughout this work. With this in mind, it is important to understand the results presented here as a statement of whether a given problem is multimodal or not and a discussion of the relative multimodality of various problems, not as an exhaustive study of all local optima present within the design space in question.

Considering all three cross-sectional degrees of freedom simultaneously, as in the 111_00 case, produces a larger increase in multimodality than a simple monotonic relationship would suggest based on the previous results. None of the reported values, save G_{m}, are unprecedented, taking values that are at most marginally different from those obtained in 101_00 or 011_00 – the next most multimodal cross-section only cases. However, a more than 4-fold increase in G_m is registered relative to those earlier cases. This indicates that while simply locating the minimum drag in this case would require approximately the same effort as the 101_00 case, if one wanted to fully explore the design space a more robust method would be necessary here than in the other cases examined thus far .

The local optima taper distributions for all relevant cross-section-only cases are shown in Figure 3. It is apparent from this figure that the behaviour of taper is linked with the other available degrees of freedom. Both 011_00 and 110_00 produce a single clearly distinct planform shape – two optima are shown for the former because variations in tip geometry have produced a roughly 1.3 drag count change in performance, but the overall planforms are highly similar. While both produce a single general planform shape, the planform in question changes substantially between cases, demonstrating a coupling between the optimal taper distribution and the ability to modify section shape and twist design variables. When all three cross-sectional degrees of freedom are considered in 111_00, the resulting design space appears to simply be the

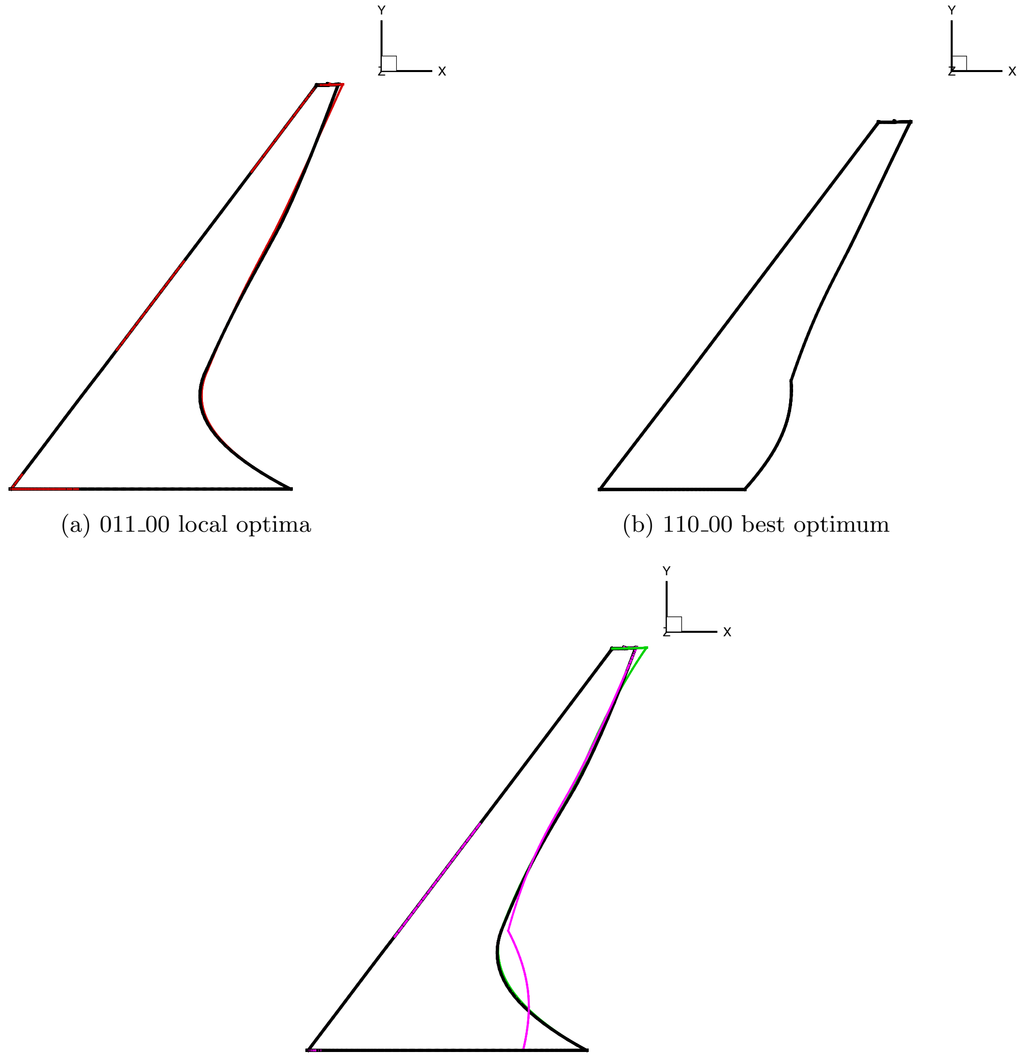

(a) 011_00 local optima

(b) 110_00 best optimum

(c) 111_00 local optima

Figure 3: Wing planforms for transonic cross-section only optimization. Best optimum in black.

superposition of the two global optima found previously. If this linear behaviour persists for taper with regards to other degrees of freedom or various constraints, then there would be a significant potential for multimodality in the taper design variables.

Spanwise twist distributions for applicable cross-section-only cases are provided in Figure 4. These results mirror and reinforce what was observed in the ADODG CRM case. The degree of multimodality in twist appears low and consistent, with each case producing between one and three twist distributions. The form of these local optima appears to be highly sensitive to the availability of other design variables, with the twist distributions varying notably between cases. The lift distributions plotted in Figure 5 reveal that all local optima find largely identical near-elliptical lift distributions, but with significantly varying total drag. As these trends are consistent across all examined transonic cases with both section and twist degrees of freedom, this strengthens the earlier conclusion that twist and section in concert permit multiple avenues to produce an elliptical lift distribution.

The overall conclusion we can draw is that in the cross-section-only optimization of a wing at transonic speed, multimodality can be present in small amounts, particularly under lightly-constrained conditions. When axial control is added, the presence of multimodality becomes far clearer than in the cross-section only cases. Both such cases – 111_01, 111_11 – exhibit clear and significant multimodality, with the degree

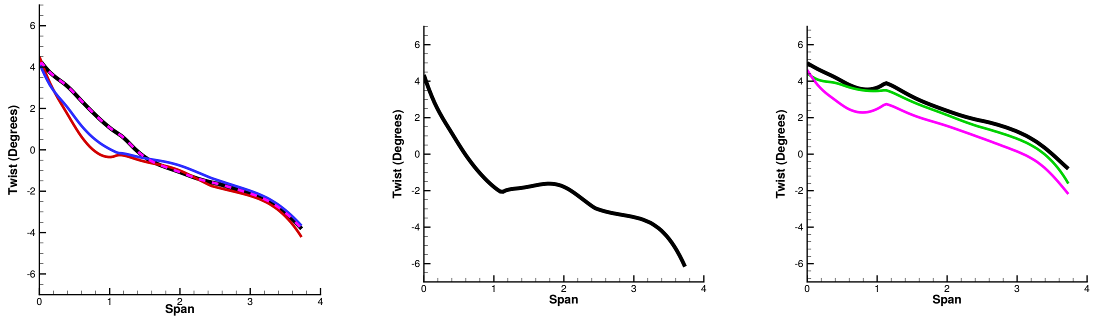

(a) 101_00 (b) 110_00 (c) 111_00

Figure 4: Twist distributions for relevant cross-section only transonic optimization local optima. Best optimum in black.

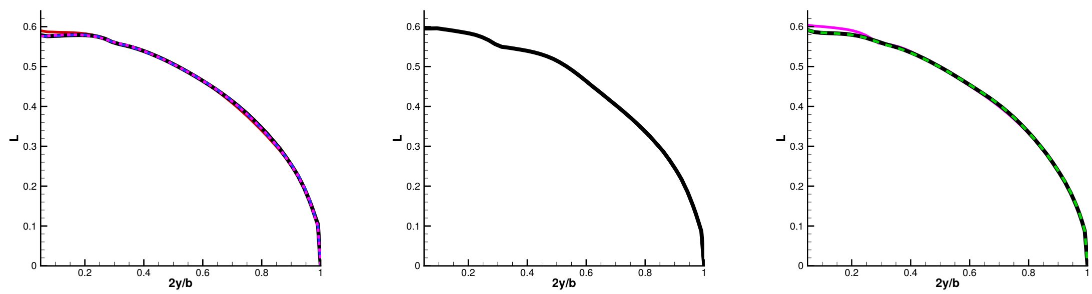

(a) 101_00 (b) 110_00 (c) 111_00

Figure 5: Spanwise lift distributions for relevant cross-section only transonic optimization local optima. Best optimum in black.

increasing as each additional degree of freedom is added. Before discussing axial results, it bears repeating that the forthcoming cases are high dimensional and lightly constrained – as such they do not represent practical wing design problems. The primary utility of these cases is to establish an upper bound on the multimodality that may be expected in a given class of problem. The question of what effect the addition of various constraints may have on mitigating this multimodality is addressed later.

The first case with axial freedom, 111_01, permitting all cross-sectional freedom as well as dihedral, produces a clear increase in multimodality. Most convincingly, a 6-fold increase in P_m and nearly 8-fold increase in G_m both indicate a design space that is considerably more demanding to navigate than the cross-section-only cases. If one examines the taper distributions in Figure 6a it is apparent that the taper multimodality exhibited in 111_00 has again vanished, with all geometries finding largely similar variations on the previously-noted concave inboard trailing edge design. The most obvious source of multimodality in this case is dihedral, where at least three dihedral modes, depicted in Figure 6b, are clear. The first two are variations on the expected dihedral-up and dihedral-down designs – though with considerable variance in the inboard dihedral distribution, ranging from relatively flat inboard sections to reflexed designs where the wing droops down before sweeping up into a hook-like dihedral. Even if one groups these variations into a single local optimum, a third mode is apparent in green, characterized by a flat inboard leading to a largely linear dihedral-down configuration. This stands in contrast to the mostly rounded dihedral designs of other minima.

Adding sweep control, as in the 111_11 case, produces still further multimodality. While neither the number of local optima nor the dominance has appreciably changed relative to 111_01, the performance range has increased nearly five-fold, the multimodality metrics have continued to climb exponentionally, and as depicted by Figures 7a and 7b, the geometric variation in the local optima has expanded. Forward- and

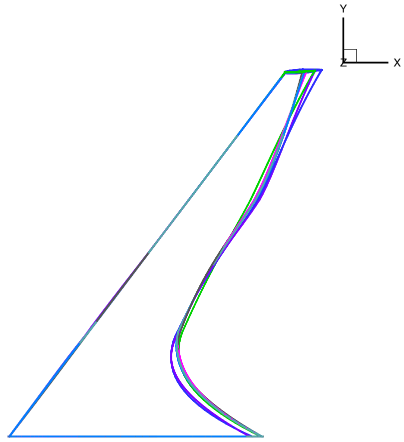

(a1) Dihedral down configurations

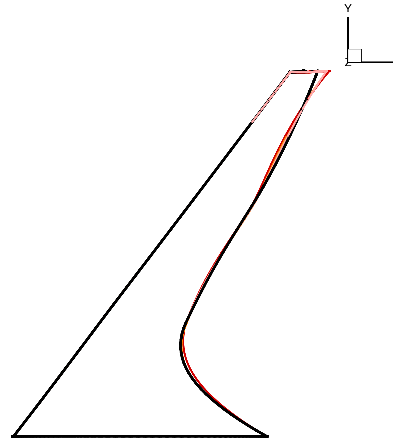

(a2) Dihedral up configurations

(a) Planform distributions

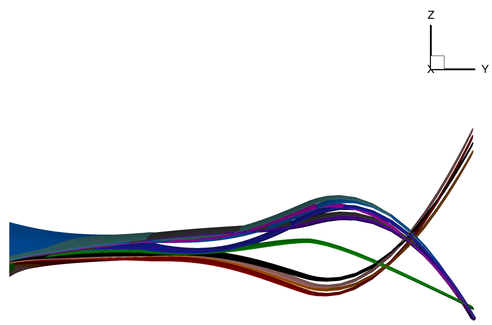

(b) Dihedral distributions

Figure 6: Local optima for transonic 111_01. Best optimum in black.

back-swept geometries are observed, and once again both positive and negative dihedral designs are well represented in the final geometries. The same trends in dihedral as 111_01 are replicated, but the geometries have become more extreme. As with previous cases, taper behaves inconsistently. Examining taper profiles, particularly inboard, multimodality is once again observed and variations on many previously-noted taper distributions are visible. The concave, convex, and linear inboard trailing edge designs all appear. As well, designs reminiscent of the global taper optimum from 110_00 have been produced.

As in the cross-section only cases, all local optima in the axial-freedom cases are shock free. The inconsistencies in the earlier noted geometric trends, or at least the difficulty in fully understanding them, may be explained in part by difficulties with convergence. The cases with axial control, particularly the 111_11 test, are high dimensional problems with extremely sparse constraints, representing a significant challenge to the optimizer – far more than the lower-dimensional problems considered previously, or the more constrained problems discussed later. This is expressed in much higher failure rates, and much shallower convergence for cases which do succeed. Therefore, we find ourselves exploring a highly multimodal design space that is also inherently difficult to traverse. Nevertheless, the multimodal nature of these problems is clear, and we can confidently conclude that the inclusion of large-scale planform deformations and non-planar configurations has a substantial impact on multimodality, particularly in lightly-constrained problems.

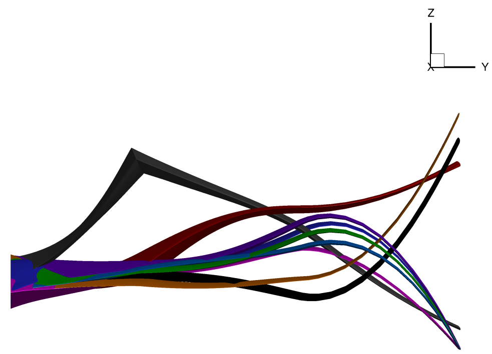

(a) Dihedral distributions

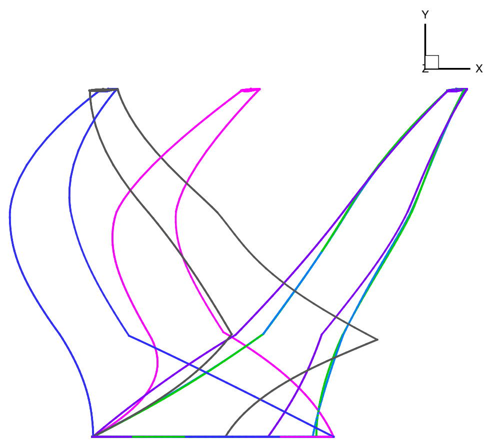

(b1) Dihedral down configurations

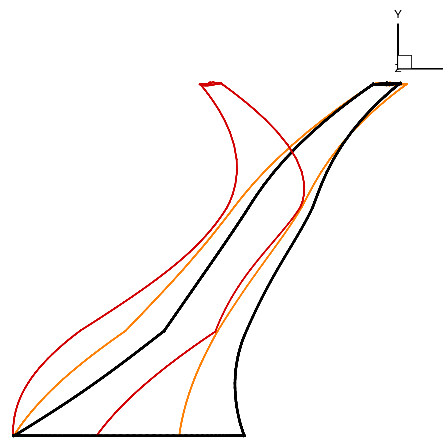

(b2) Dihedral up configurations

(b) Planform distributions

Figure 7: Local optima for transonic 111_11. Best optimum in black.

F. Subsonic Degree of Freedom Study

The second set of cases seek examine the impact of flow regime on multimodality. This is accomplished by repeating the degree of freedom study undertaken in Section E but at a subsonic Mach number of 0.5, and a Reynolds number of 2.94 million. Other than this variation in flow regime, the problem definition is identical to that of the transonic degree of freedom tests, consisting of the baseline viscous lift-constrained drag minimization with minimum volume and fixed projected area. Once again, apart from generous upper and lower bounds on each design variable, no additional constraints are used. The naming convention from those previous tests is duplicated here as well, with each test case denoted by a five digit code of the form xxx_xx, where the first three digits refer to the cross sectional degrees of freedom (twist, taper, section control), the latter two correspond to the sweep and dihedral axial degrees of freedom, and each digit takes a 0 or 1 to denote whether that degree of freedom is permitted in a particular test.

Numerical results from each of the seven subsonic cases are tabulated in Table 4. Overall, it is observed that the general trends from the transonic cases are consistent across the subsonic cases as well. Regarding the cross-section only tests – 001_00 through 111_00 – the immediately apparent effect of moving to a subsonic regime is a clear reduction in the degree of observed multimodality. While varying, but consistently small, amounts of multimodality were noted in the transonic cross-section only cases, here all except the 111_00 case are fully unimodal. As in the transonic case, the 111_00 is the most multimodal of the cross-sectional cases

Table 4: Degree of Freedom Multimodality Study Numerical Results. M=0.50, Re=2.94 million. Each test is identified using a five digit code of the format xxx_xx. From left to right each digit indicates whether twist, taper, section shape, sweep, or dihedral freedom is permitted.

Test Code	Successful Tests	Optima	Best (drag counts)	Range (drag counts)	Dominance	P_m	G_m
001_00	33	1	193	0.0	100%	0.00	0.00
011_00	33	1	192	0.0	100%	0.00	0.00
101_00	33	1	193	0.0	100%	0.00	0.00
110_00	33	1	201	0.0	100%	0.00	0.00
111_00	17	2	191	1.9	94%	0.06	4.47
111_01	10	7	183	14.9	30%	3.72	52.72
111_11	9	8	180	7.3	22%	2.08	52.08

– though here that is merely by virtue of being the only cross-sectional case to exhibit any multimodality whatsoever.

Turning to geometry, to parallel the transonic discussion we will primarily focus on the taper and twist design variables in the cross-section-only discussion. What we find is quite consistent with the transonic regime. Figure 8 displays the 011_00, 110_00, and 111_00 taper distributions, showing the same patterns as observed in the transonic cases. Once again, both 011_00 and 110_00 produce one dominant taper distribution, as before the 011_00 case generates a concave inboard trailing edge and tapered wingtip, while the 110_00 case prefers a more convex inboard – here almost linear with a slight bulge – leading to a wingtip similar to that in the 011_00 case. Finally, 111_00 again produces two planforms which correspond to each of the planforms found in the other two cases. While the geometries do vary – for instance the concave subsonic 111_00 local optimum and the transonic 110_00 best optimum both bulge at the inboard trailing edge more than the subsonic 110_00 best optimum, the trends are clearly consistent.

Regarding twist, relevant distributions are found in Figure 9. While the 101_00 and 110_00 twist distributions bear little resemblance to their transonic counterparts, the form of the 111_00 distributions is remarkably similar across both flow regimes. Based on this evidence we can conclude that in the subsonic regime, twist appears to be either unimodal or possessing a small number of local optima. Figure 10 illustrates that all subsonic section-only local optima are able to locate an elliptical lift distribution.

Examing the cases with axial degrees of freedom permitted – 111_01, 111_11 – the trends become less clear. The initial jump in multimodality between 111_00 and 111_01, which reflects the effect of adding dihedral control, is larger in the subsonic case than the transonic. However, the subsequent growth of multimodality is less steep, with the peak subsonic P_m value being approximately 30% the peak transonic value, and G_m only achieving a value approximately two-thirds the magnitude of the maximum in the transonic tests. P_m also no longer monotonically increases as control is added, dropping between the subsonic 111_01 and 111_11 tests while the same step caused an approximately 400% increase in P_m in the transonic regime. The relative multimodality for cases between flow regimes is also inconsistent: measured by the P_m and G_m values, the 111_11 case produces significantly more multimodality in the transonic than subsonic regime, while the 111_01 case shows the opposite trend, with the subsonic metric values being notably higher than the transonic. In all cases, the relatively small variations in the number of optima and dominance values are deceptive, with variations in P_m and G_m being predominantly driven by changes to the range and distribution of the local optima. As in the transonic examinations, slow convergence and flow solve or mesh movement failures continue to be problematic for the subsonic axial cases, particularly the 111_11 test, though less so than in the transonic results.

Overall, a decrease in multimodality is noted when moving to the subsonic region, as the majority of examined cases showed significant reductions in all relevant multimodality metrics, the one notable exception being 111_01.

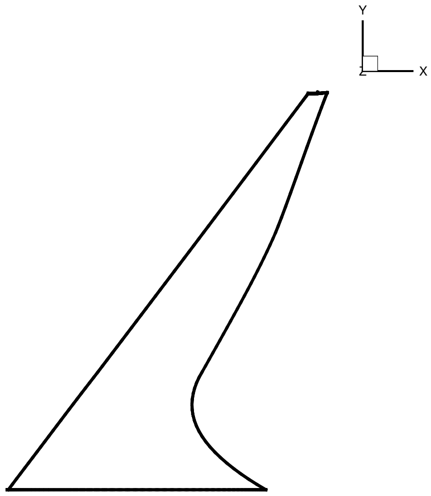

(a) 011_00 best optimum

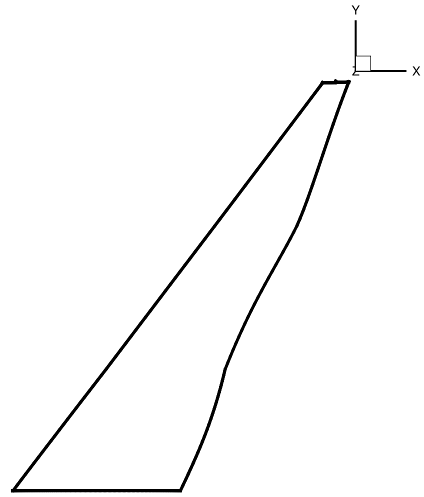

(b) 110_00 best optimum

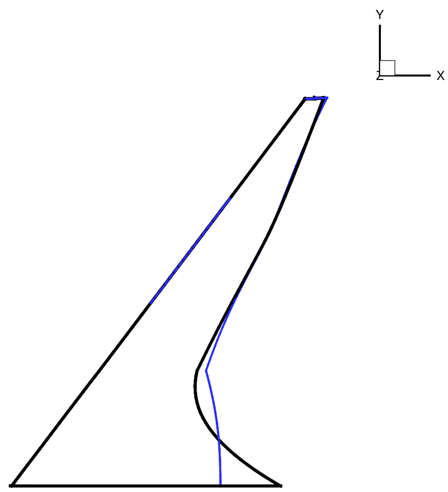

(c) 111_00 local optima

Figure 8: Wing planforms for subsonic cross-section only optimization. Best optimum in black.

G. Constraint Study

The previous tests examined the relationship that degrees of freedom and flow regime have with multimodality. While the minimally-constrained nature of the previous cases are useful for establishing an upper bound on the multimodality that may be present in a problem, understanding the presence of multimodality in more practical problems requires a focused study of the relationship between constraints and multimodality. The ultimate goal of this study is to determine whether multimodality is present in a high dimensional problem subject to practical geometric constraints. The same baseline problem is considered as in the degree of freedom test; however now rather than varying the degrees of freedom as in earlier tests, the degrees of freedom are fixed and the constraints are varied. The transonic 111_11 case is used as an "unconstrained" baseline; this was selected as a high dimensional case permits us to examine a wide variety of constraints. If common structural or manufacturability constraints – such as a minimum thickness or straight leading and trailing edges – can eliminate multimodality in a case such as this, then it is highly likely that most or all practical problems of lower dimensionality are unimodal. However, if multimodality persists even under tightly constrained circumstances, then multimodality must be accounted for even in practical optimization problems.

A selection of common constraints are examined in different combinations, in a similar manner to the

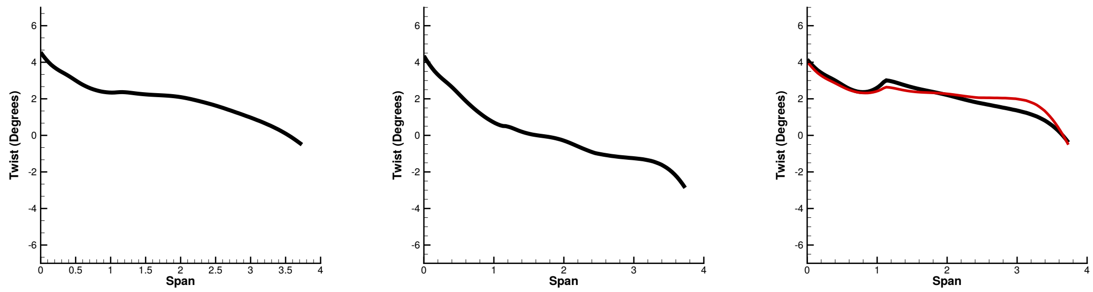

(a) 101_00 (b) 110_00 (c) 111_00

Figure 9: Twist distributions for relevant cross-section only subsonic optimization local optima. Best optimum in black.

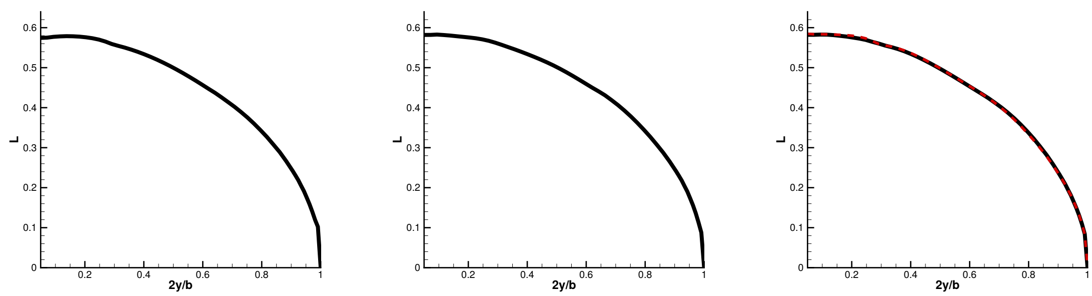

(a) 101_00 (b) 110_00 (c) 111_00

Figure 10: Lift distributions for relevant cross-section only subsonic optimization local optima. Best optimum in black.

above degree of freedom investigation. The constraints examined are linear twist and taper – requiring the design variables to follow a linear distribution between the root and tip of each FFD volume, though not necessarily across volumes – minimum thickness – constraining thickness variations to between 0.85 and 1.5 times the initial thickness of each section – constant sweep and dihedral – which enforce constant sweep or dihedral angles across both FFD volumes – and pitching moment – which constrains the pitching moment coefficient to be greater than or equal to -0.17. Each test is denoted by a five digit code of the form xxx_xx where each digit from left to right corresponds to the twist, taper, section control, sweep, and dihedral. If a constraint is present on a degree of freedom, that digit will take one of several letters to indicate the kind of constraint used, corresponding to (L)inear, (M)inimum or (C)onstant. The suffix "NPM" is appended to certain names to indicate that a pitching moment constraint is not enforced.

Fundamentally, this test was envisioned as starting with a clearly multimodal, unconstrained problem, and gradually adding constraints to determine when multimodality is reduced or eliminated entirely. The most practical optimizations require fully straight leading and trailing edges, a common requirement in design problems achieved through a combination of linear constraints. Straight leading and trailing edges in the xy plane require linear taper and constant sweep – denoted as 0L0_C0. Straight leading and trailing edges in the yz plane require constant dihedral – 000_0C. Tests with all three of these constraints will have fully straight leading and trailing edges.

One of the degrees of freedom which displayed the most varied behaviour in the previous sections was taper, and as such the linear taper constraint was studied first; further tests were selected by adding one constraint at a time to each previous test. Numerical results from these tests are found in Table 5, for clarity, the cases are grouped according to major shared properties, specifically: lack of a pitching moment constraint, lack of any leading and trailing edge restrictions, partial leading and trailing edge restrictions, or

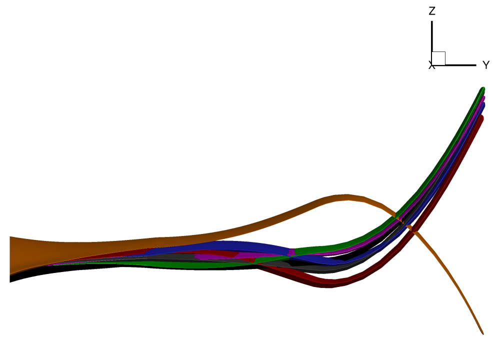

(a) Dihedral distributions

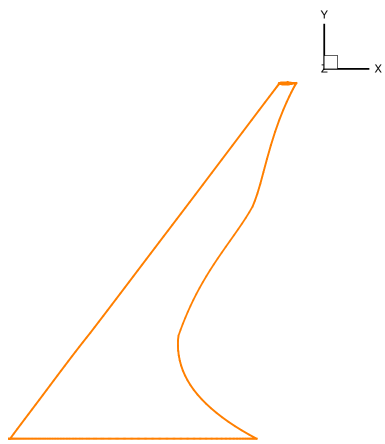

(b1) Dihedral down configurations

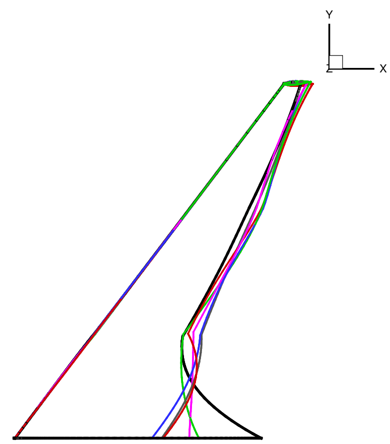

(b2) Dihedral up configurations

(b) Planform distributions

Figure 11: Local optima for subsonic 111_01. Best optimum in black.

fully straight leading and trailing edges.

Multimodality, measured by any of the reported values, tends to decrease as constraints are added – though this trend is not monotonic. More interesting is the persistence of multimodality in tightly constrained problems, with every examined case indicating multimodality of some degree. This includes the most stringently constrained tests, those requiring fully straight leading and trailing edges. This is accomplished in cases LL0_CC and 0L0_CC by requiring linear taper, constant dihedral, and constant sweep – in addition to linear twist in the former case. The final optima from the LL0_CC case are plotted in Figures 13a and 13b. Comparing these to the optima produced by the unconstrained, transonic 111_11 case which provides the basis of this test – shown in Figures 7a and 7b – it is apparent that multimodality has been reduced in the constrained case. This conclusion is corroborated by the reduction in nearly every reported value in Table 5 between these two cases; in particular the dihedral-up mode has disappeared and dihedral appears largely unimodal with perhaps some variation in angle between modes. However, there remains clear multimodality within the planform shapes depicted in Figure 13b. Beyond the forward and back-swept geometries, there is evidence of at least two, perhaps three, distinct taper distributions within these optima. The only multimodality metric which does not drop in the constrained case is G_m, which increases significantly relative to the unconstrained case. This is driven by the same trend driving the P_m value down: the relative rarity of outlying optima in this design space. While the overall geometric spread of the optima has clearly decreased relative to the unconstrained test, the relative dominance of the more similar optima

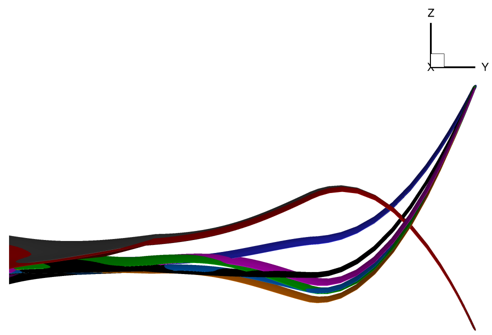

(a) Dihedral distributions

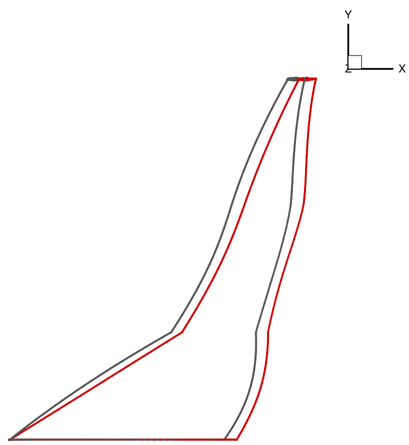

(b1) Dihedral down configurations

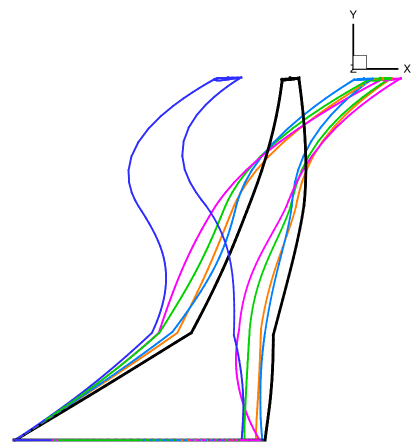

(b2) Dihedral up configurations

(b) Planform distributions

Figure 12: Local optima for subsonic 111_11. Best optimum in black.

has increased, rendering those outlying optima more difficult to locate. This is a recurring trend in the constrained test results.

Removing the linear twist constraint, the 0L0_CC case, produces significant changes to the design space. Table 5 shows that doing so causes the number of optima to double and the dominance to fall by a third, suggesting a design space that is appreciably more multimodal; this is supported by Figure 14 which shows a notable increase in the variety of planform designs and comparable variability in dihedral relative to LL0_CC. However, examining the P_m and G_m values reveals a more nuanced perspective. Relative to the linear-twist constrained case, the performance range has dropped and the optima have clustered more closely around the minimum drag; this produces a marginally reduced P_m value, suggesting that it is somewhat easier to get close to the minimum drag in this case, despite the additional optima. However, most of the geometrically outlying optima are relatively rare, many with only a single example out of 15 successful initial geometries. This produces a large G_m and indicates that while one may be able to use a relatively small number of samples and have confidence that one is not losing significant performance in the final design, ensuring that one has located all major geometric modes requires a far more in-depth study of the design space.

The absence of any dihedral-up designs in either of the two constrained tests examined thus far is somewhat unexpected as there is little reason to suspect that the variety of dihedral-up modes observed in other cases would be entirely eliminated by a constant dihedral constraint. In fact, a constant-dihedral-up optimum is located by 000_0C, shown in Figure 15, which enforces only a constant dihedral constraint. This

Table 5: Constrained numerical results. M=0.85, Re=5 million. 000_00_NPM case corresponds to transonic 111_11. "L" requires a design variable to have a linear distribution, "C" requires a constant distribution, "M" enforces a strict minimum value, and "NPM" denotes that a pitching moment constraint is not enforced.

Test Code	Successful Tests	Optima	Best (drag counts)	Range (drag counts)	Dominance	P_m	G_m
No Pitching Moment Constraint							
000_00_NPM	10	9	178	34.8	20%	6.10	74.41
0L0_00_NPM	10	10	178	20.4	10%	2.56	45.75
0L0_0C_NPM	15	4	183	11.0	67%	0.50	6.84
No LE/TE Restrictions							
0L0_00	10	10	183	11.2	10%	3.56	21.58
0LM_00	10	10	189	7.0	10%	1.43	17.83
LL0_00	10	9	185	20.8	20%	3.81	51.60
LLM_00	12	10	189	15.7	17%	1.86	36.01
Some LE/TE Restrictions							
000_0C	10	6	186	6.9	50%	0.92	24.68
0L0_0C	10	4	187	7.1	70%	0.54	6.84
0L0_C0	12	4	183	4.4	42%	0.91	104.73
0LM_0C	13	4	194	15.4	77%	0.97	59.60
0LM_C0	12	6	191	14.8	33%	1.16	60.16
LL0_0C	13	4	187	4.5	54%	0.26	44.04
LL0_C0	15	9	183	7.4	27%	1.36	40.87
Straight LE/TE							
0L0_CC	15	8	188	7.3	40%	0.98	84.01
LL0_CC	15	4	188	8.9	60%	1.13	129.13

particular optimum was located by just one of the 17 attempted geometries, presenting the possibility that if such modes do exist within the 0L0_CC and LL0_CC design spaces they are non-dominant and potentially difficult to locate.

This result also highlights the impact that constant sweep and dihedral constraints have on multimodality, even when applied in isolation. Throughout the tested cases, the addition of a constant sweep or dihedral constraint – denoted as xxx_xC or xxx_Cx – has a consistent negative impact on multimodality, indicated by large reductions in the number of located optima and P_m and corresponding increases in dominance. This can be seen by comparing any of the reported values for 0L0_00 with 0L0_0C or 0L0_C0, the results for 0LM_00 with its sweep or dihedral constrained varients, or any other relevant cases. The one consistent exception to this relationship is G_m, which quite frequently is observed to increase with the addition of these constraints. This is related to the previous discussion on the possible presence of dihedral-up modes in constant-dihedral cases, suggesting that these constraints may not eliminate all geometrically outlying modes, but merely render some much rarer and therefore harder to locate. This should serve to caution readers against assuming that merely by constraining critical design variables one can ensure an easy-to-navigate design space.

Among the other tested constraints – linear twist, linear taper, minimum thickness, and pitching moment – few strong, consistent relationships are apparent. Pitching moment appears to have at most a moderate impact on multimodality, serving to flatten performance variations and redistribute local optima relative to an unconstrained test but having little to no effect on the number of local optima found, and a highly inconsistent relationship with the resulting P_m and G_m values. A similar relationship is noted for the

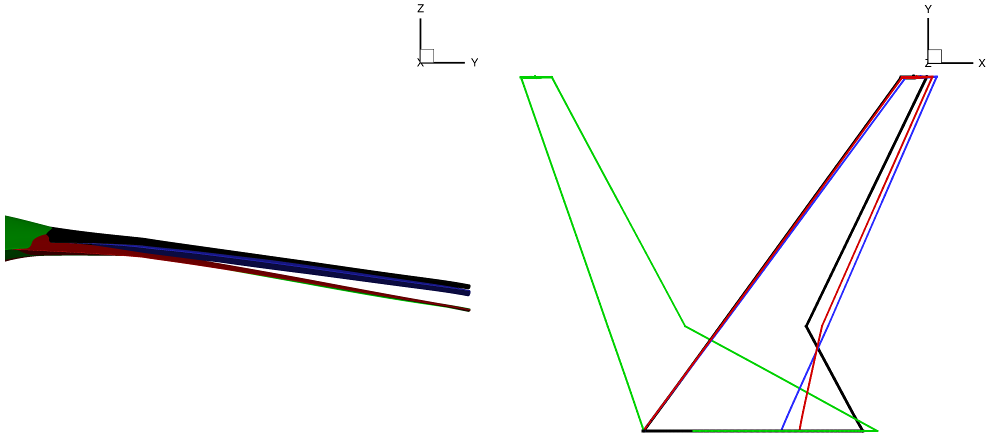

(a) Dihedral distributions (b) Wing planforms

Figure 13: Local optima for LL0_CC. Best optimum in black.

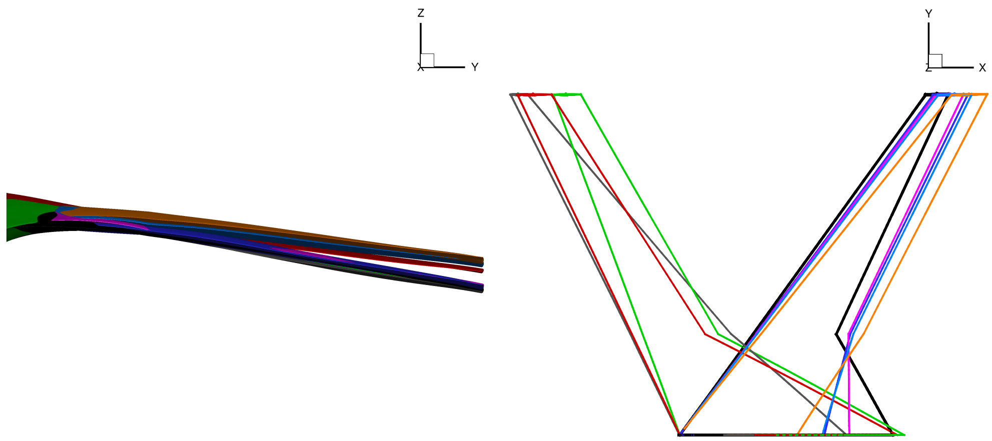

(a) Dihedral distributions (b) Wing planforms

Figure 14: Local optima for 0L0_CC. Best optimum in black.

minimum thickness constraint, whose impact is largely confined to a large reduction in the performance range achieved by the optimizer. The impact of linear taper and linear twist appear to be highly related to other problem parameters, including which other constraints are applied simultaneously. As a result, both constraints have highly variable relationships, appearing to cause large increases in various metrics at some times, and large decreases in others. However, it appears that linear taper, at least, has a minimal impact when enforced in isolation. Comparing the numerical results in Table 5 for 000_00_NPM and 0L0_00_NPM one can see that while the performance range has decreased there is little notable change in the number of optima or dominance; indeed examining the local optima produced by the 0L0_00_NPM case in Figures 16 shows little material change in the optima obtained relative to those in in Figures 7a and 7b. That considered, while the number of local optima has not changed, large reductions in both P_m and G_m point to a redistribution of those optima within the design space, producing a problem that is significantly less burdensome to navigate.

Overall, the most important conclusion that can be drawn from these results is the consistent presence of multimodality in nearly every examined design space. This demonstrates that multimodality can and does exist in well-constrained practical wing optimization problems. The data presented here strongly suggest that using a gradient-based optimization algorithm starting from a single initial geometry will not be sufficient to ensure that the global optimum is found.

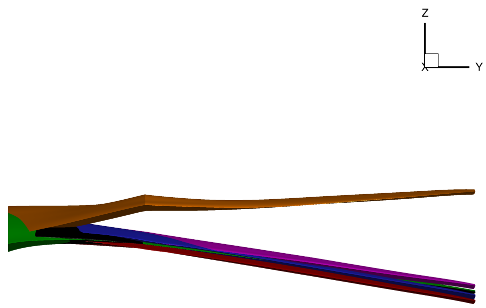

Figure 15: Dihedral distributions for 000_0C local optima

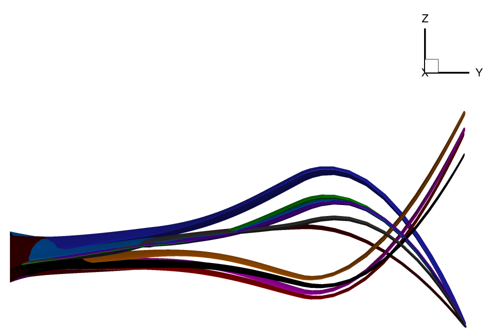

(a) Dihedral distributions

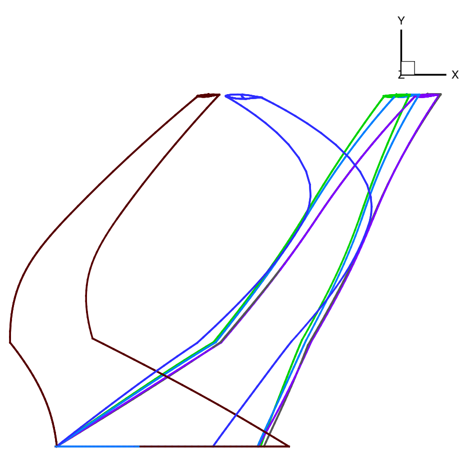

(b1) Dihedral down configurations

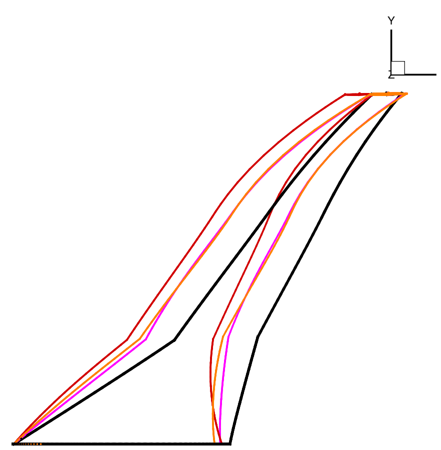

(b2) Dihedral up configurations

(b) Planform distributions

Figure 16: Local optima for 0L0_00_NPM. Best optimum in black.

IV. Conclusion

This paper has presented a thorough parametric study of multimodality in the RANS-based optimization of a wing. The data indicate that for wings in both subsonic and transonic flows, optimization of the cross-section through twist, taper, and section shape is somewhat multimodal, while permitting changes to the sweep and dihedral produces moderately to highly multimodal design spaces. Generally, wing optimization in subsonic flows appears less multimodal than in transonic flows, but clear evidence for multimodality is present in both. Multimodality in cases permitting large planform deformations can be reduced but not eliminated by enforcing a requirement for straight leading and trailing edges. Other constraints, such as linear twist or taper, minimum thickness, or a pitching moment constraint appear to have inconsistent impacts on multimodality; however multimodality remains to some degree under all examined constraint combinations

Overall, multimodality has been shown to be an intrinsic component of the design space for wing optimization under many circumstances. The multimodality observed is of variable degree, but appears persistent across various flow regimes and constraints. Based on these results, while global optimization algorithms may not be necessary for all problems, if one chooses to use a local optimization scheme for the aerodynamic shape optimization of a wing, the possible presence of multimodality must be understood as an inherent risk.

References

[1] IATA, "IATA Annual Review 2017," 2017.

[2] Jameson, A., "Aerodynamic design via control theory," *Journal of Scientific Computing*, Vol. 3, No. 3, 1988, pp. 233–260.

[3] Holst, T. and Pulliam, T. H., "Aerodynamic shape optimization using a real-number-encoded genetic algorithm," Tech. Rep. 2001-2473, AIAA, June 2001.

[4] Zingg, D. W., Nemec, M., and Pulliam, T. H., "A comparative evaluation of genetic and gradient-based algorithms applied to aerodynamic optimization," *European Journal of Computational Mechanics/Revue Européenne de Mécanique Numérique*, Vol. 17, No. 1-2, 2008, pp. 103–126.

[5] Sefrioui, M., Srinivas, K., and Periaux, J., "Aerodynamic Shape Optimization Using a Hierarchical Genetic Algorithm," *European Congress on Computational Methods in Applied Sciences and Engineering*, Barcelona, 2000, pp. 20–38.

[6] Husain, A., Lee, K., and Kim, K., "Enhanced multi-objective optimization of a dimpled channel through evolutionary algorithms and multiple surrogate methods," *International Journal for Numerical Methods in Fluids*, Vol. 66, No. 6, 2011, pp. 742–759.

[7] Lombardi, A., Ferrari, D., and Santos, L., "Aircraft Air Inlet Design Optimization via Surrogate-Assisted Evolutionary Computation," *Evolutionary Multi-Criterion Optimization*, Springer, 2015, pp. 313–327.

[8] Shahrokhi, A. and Jahangirian, A., "A surrogate assisted evolutionary optimization method with application to the transonic airfoil design," *Engineering Optimization*, Vol. 42, No. 6, 2010, pp. 497–515.

[9] Poloni, C. and Mosetti, G., "Aerodynamic Shape Optimization by Means of Hybrid Genetic Algorithm," *Zeitschrift fuer Angewandte Mathematik und Mechanik*, Vol. 76, No. S3, 1996, pp. 247–250.

[10] Sóbester, A. and Keane, A., "Empirical comparison of gradient-based methods on an engine-inlet shape optimization problem," *9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, No. 2002-5507, 2002.

[11] Chernukhin, O. and Zingg, D. W., "Multimodality and Global Optimization in Aerodynamic Design," *AIAA Journal*, Vol. 51, No. 6, 2013, pp. 25–34.

[12] Du, X., Amrit, A., Thelen, A., Leifsson, L., Zhang, Y., Han, Z. H., and Koziel, S., "Aerodynamic Design of a Rectangular Wing in Subsonic Inviscid Flow by Direct and Surrogate-based Optimization," *35th AIAA Applied Aerodynamics Conference*, No. 2017-4366, 2017.

[13] Poole, D., Allen, C., and Rendall, T., "Global Optimization of Multimodal Aerodynamic Optimization Benchmark Case," *35th AIAA Applied Aerodynamics Conference*, No. 2017-4365, 2017.

[14] Bons, N. P., He, X., Mader, C. A., and Martins, J. R. R. A., "Multimodality in Aerodynamic Wing Design Optimization," *35th AIAA Applied Aerodynamics Conference*, No. 2017-3753, 2017.

[15] Streuber, G. M. and Zingg, D. W., "Investigation of multimodality in aerodynamic shape optimization based on the Reynolds-Averaged Navier-Stokes Equations," *35th AIAA Applied Aerodynamics Conference*, No. 2017-3752, 2017.

[16] Koo, D. and Zingg, D. W., "Investigation into Aerodynamic Shape Optimization of Planar and Nonplanar Wings," *AIAA Journal*, Vol. 56, No. 1, 2018, pp. 1–14.

[17] Yu, Y., Lyu, Z., Xu, Z., and Martins, J., "On the influence of optimization algorithm and initial design on wing aerodynamic shape optimization," *Aerospace Science and Technology*, Vol. 75, 2018, pp. 183–199.

[18] Lee, C., Koo, D., Telidetzki, K., Buckley, H., Gagnon, H., and Zingg, D. W., "Aerodynamic shape optimization of benchmark problems using Jetstream," *53rd AIAA Aerospace Sciences Meeting*, No. 2015-0262, 2015.

[19] Shi-Dong, D., Chen, C., and Nadarajah, S., "Adjoint-Based Aerodynamic Optimization of Benchmark CRM Wing," *35th AIAA Applied Aerodynamics Conference*, No. 2017-3755, 2017.

[20] Lyu, Z., Kenway, G., and Martins, J., "Aerodynamic shape optimization investigations of the common research model wing benchmark," *AIAA Journal*, Vol. 53, No. 4, 2015, pp. 968–985.

[21] Poole, D., Allen, C., and Rendall, T., "Global Optimization of Wing Aerodynamic Optimization Case Exhibiting Multimodality," *Journal of Aircraft*, 2018, pp. 1–16.

[22]Sederberg, T. W. and Parry, S. R., “Free-Form Deformation of Solid Geometric Models,” *13th Annual Conference on Computer Graphics and Interactive Techniques*, No. 4, Dallas, Texas, 1986, pp. 151–160.

[23]Joe, S. and Kuo, F. Y., “Remark on algorithm 659: Implementing Sobol's quasirandom sequence generator,” *ACM Transactions on Mathematical Software (TOMS)*, Vol. 29, No. 1, 2003, pp. 49–57.

[24]Antonov, I. A. and Saleev, V., “An economic method of computing LP τ-sequences,” *USSR Computational Mathematics and Mathematical Physics*, Vol. 19, No. 1, 1979, pp. 252–256.

[25]Gyorgy, A. and Kocsis, L., “Efficient Multi-Start Strategies for Local Search Algorithms,” *Journal of Artificial Intelligence Research*, Vol. 41, 2011, pp. 407–444.

[26]Ugray, Z., Lasdon, L., Plummer, J., Glover, F., Kelly, J., and Marti, R., “A Multistart Scatter Search Heuristic for Smooth NLP and MINLP Problems,” *Metaheuristic Optimization via Memory and Evolution*, Vol. 30, 2005, pp. 25–57.

[27]Tu, W. and Mayne, R., “Studies of Muilti-Start Clustering for Global Optimization,” *International Journal for Numerical Methods in Engineering*, Vol. 53, 2002, pp. 2239–2252.

[28]Tu, W. and Mayne, R., “An Approach to Multi-Start Clustering for Global Optimization with Non-Linear Constraints,” *International Journal for Numerical Methods in Engineering*, Vol. 53, 2002, pp. 2253–2269.

[29]Schwabacher, M. and Gelsey, A., “Multilevel Simulation and Numerical Optimization of Complex Engineering Designs,” *Journal of Aircraft*, Vol. 35, No. 3, 1998, pp. 387–397.

[30]Gagnon, H. and Zingg, D. W., “Two-Level Free-Form and Axial Deformation for Exploratory Aerodynamic Shape Optimization,” *AIAA Journal*, Vol. 53, No. 7, 2015, pp. 2015–2026.

[31]Hicken, J. E. and Zingg, D. W., “Aerodynamic optimization algorithm with integrated geometry parameterization and mesh movement,” *AIAA Journal*, Vol. 48, No. 2, 2010, pp. 400–413.

[32]Osusky, L., Buckley, H., Reist, T., and Zingg, D. W., “Drag minimization based on the Navier–Stokes equations using a Newton–Krylov approach,” *AIAA Journal*, Vol. 53, No. 6, 2015, pp. 1555–1577.

[33]Hicken, J. E. and Zingg, D. W., “Parallel Newton-Krylov solver for the Euler equations discretized using simultaneous approximation terms,” *AIAA Journal*, Vol. 46, No. 11, 2008, pp. 2773–2786.

[34]Osusky, M. and Zingg, D. W., “Parallel Newton–Krylov–Schur Flow Solver for the Navier–Stokes Equations,” *AIAA Journal*, Vol. 51, No. 12, 2013, pp. 2833–2851.

[35]Gill, P. E., Murray, W., and Saunders, M. A., “SNOPT: An SQP algorithm for large-scale constrained optimization,” *SIAM Review*, Vol. 47, No. 1, 2005, pp. 99–131.

[36]Spalart, P. and Allmaras, S., “A one-equation turbulence model for aerodynamic flows,” *La Recherche Aerospatiale*, Vol. 1, 1994, pp. 5–21.

[37]Nocedal, J. and Wright, S., *Numerical Optimization*, Vol. 35, Springer Science, 1999.

[38]Hicken, J. E. and Zingg, D. W., “A simplified and flexible variant of GCROT for solving nonsymmetric linear systems,” *SIAM Journal on Scientific Computing*, Vol. 32, No. 3, 2010, pp. 1672–1694.

[39]De Sturler, E., “Truncation strategies for optimal Krylov subspace methods,” *SIAM Journal on Numerical Analysis*, Vol. 36, No. 3, 1999, pp. 864–889.

[40]Levy, D. W., Laflin, K. R., Tinoco, E. N., Vassberg, J. C., Mani, M., Rider, B., Rumsey, C., Wahls, R. A., Morrison, J. H., Brodersen, O. P., Crippa, S., Mavriplis, D., and Murayama, M., “Summary of data from the fifth AIAA CFD drag prediction workshop,” *51st AIAA Aerospace Sciences Meeting*, No. 2013-0046, 2013.