

On the Risk of Local Optima in Aerodynamic Shape Optimization

Gregg M. Streuber* and David. W. Zingg†

Institute for Aerospace Studies, University of Toronto, Toronto, Ontario, M3H 5T6, Canada

A gradient-based multistart method based on a set of seventeen to thirty-three random initial geometries is used to examine the risk associated with multimodality when applying gradient-based optimization to aerodynamic shape optimization. Aerodynamic shape optimization problems typical of detailed, preliminary, and exploratory design are shown to consistently present design spaces with multiple local optima. In the case of detailed design, the risk of converging to a local optimum with performance significantly inferior to that of the best local optimum found is reduced due to the ability of a well-designed initial geometry, which is often available for such problems, to converge to a well-performing local optimum. In problems permitting increased geometric freedom typical of preliminary design the risk associated with multimodality is much higher. This risk is further exacerbated in exploratory cases where high geometric freedom is combined with limited knowledge of the design space in question and hence greater differences between available initial geometries and the optimal geometry. Therefore, for preliminary and exploratory design, allocating resources toward addressing multimodality can significantly reduce the risk of overlooking a superior optimum.

I. Introduction

Aerodynamic shape optimization is a field coupling numerical optimization with computational fluid dynamics (CFD), where the output of a CFD simulation around an aerodynamic body is optimized with respect to the shape of the body. With the adoption of the discrete adjoint method [1], gradients can be computed at a cost almost independent of the number of design variables. Combined with the rise of affordable computing, this has made aerodynamic shape optimization practical for the optimization of full aircraft designs in high fidelity and has led to its widespread adoption in industry and academia for the analysis and improvement of aircraft [2–11]. Aerodynamic shape optimization problems can often be associated with two distinct phases of the design process: detailed design and preliminary design. Detailed design is performed later in the design process when a reasonable design has already been developed and further geometric changes are limited. Preliminary design is performed earlier in the design process and typically involves

*Ph.D. Candidate, gregg.streuber@utoronto.ca

†University of Toronto Distinguished Professor of Computational Aerodynamics and Sustainable Aviation, Director, Centre for Research in Sustainable Aviation, and Associate Fellow AIAA, dwz@oddjob.utias.utoronto.ca

greater geometric flexibility. Exploratory design may be considered as a third class of problem or as a subcategory within preliminary design, combining large geometric flexibility with often unknown design spaces, and searching for novel designs. Common examples of exploratory optimization include the study of hybrid wing body (HWB) geometries [3, 12], box-wing aircraft [13] or other unconventional designs. In preliminary or exploratory design, aerodynamic shape optimization is often combined with other disciplines as in multidisciplinary optimization. In the development of novel unconventional aircraft, where existing design experience is minimal or nonexistent and hence the optimal shape is likely to bear little resemblance to the initial geometry, aerodynamic shape optimization can be particularly valuable.

Gradient-based optimization methods based on the adjoint method have become widely used for both classes of problem. However, a major concern arises when a gradient-based optimization method is applied to a problem that may not be convex and therefore multiple local optima can exist. A gradient-based method will converge to a particular local optimum with no indication of whether multiple local optima are present in the design space. Global optimization methods are often gradient-free and are capable of locating global optima; such methods have been applied to aerodynamic shape optimization [14]. However, global optimization is significantly more expensive than gradient-based approaches [15], cannot guarantee convergence to the global optimum, may experience issues with repeatability, and there is little consensus in the literature on best methods for global optimization. For problems with a moderate number of local optima, methods such as Gradient-Based Multistart (GBMS) can provide an effective balance between the speed of gradient-based optimization and the robustness of global optimization methods [16]. However, in a practical setting, limits on available computing resources can restrict the use of GBMS. Therefore it is important for the designer to have some understanding of how the risk associated with the use of a local minimizer depends on the nature of the optimization problem.

Chernukhin & Zingg [16] were among the first to study multimodality - the presence of multiple local optima - in aerodynamic shape optimization, finding that two-dimensional airfoil optimization under transonic cruise conditions, subject to the Reynolds-Averaged Navier-Stokes (RANS) equations, is unimodal - presenting only a single optimum - but finding multiple local optima in both subsonic and transonic inviscid 3D wing optimizations, as well as in a highly exploratory Hybrid Wing Body (HWB) optimization in transonic, inviscid flow. Limited studies of multimodality in the Aerodynamic Design and Optimization Discussion Group (ADODG) case 5, the viscous twist and section optimization of the Common Research Model (CRM) wing-only geometry, have been undertaken by Lyu et al. [17], and Koo and Zingg [18]. Lyu et al. found evidence of multimodality in 3D optimization; however, their study observed minimal variations in geometry and performance between local optima. Koo and Zingg, examining the same problem, found it to be unimodal. With the publication of ADODG Case 6, a definitively multimodal 3D wing optimization case, a large number of investigations soon followed. The problem was originally found to be multimodal by Streuber and Zingg [19], which was confirmed in work by Poole et al. [20, 21], who studied several variations of ADODG case 6 using gradient-free global optimization methods. They located consistent evidence of multimodality, including in

relatively low-dimensional design spaces such as chord optimization. They concluded that the addition of section control or the permitting of large-scale planform changes like sweep or dihedral can dramatically increase multimodality, while some of their results suggest that the relationship between the number of local optima and the number of design variables can be complex, subject to potentially significant coupling. Yu et al. [22] investigated multimodality in the transonic, viscous optimization of the section shape and twist of the CRM wing-only geometry and found it to be slightly multimodal. Most recently, Bons et al.[23] have undertaken a study of the impact of varying the degrees of freedom and monotonic or linear chord constraints on multimodality in subsonic aerodynamic shape optimization, based on variations of ADODG Case 6 under both viscous and inviscid conditions, concluding that multimodality is minimal in subsonic optimization, and correctable with improved physics modelling and increased constraints.

The primary question of interest to a practitioner of aerodynamic shape optimization is whether or not multiple local optima are present; convex design spaces may be efficiently explored using a purely gradient-based optimization algorithm, cases with more local optima are ideal for the GBMS method of Chernukhin and Zingg [16], while for design spaces with extreme multimodality it may be necessary to incur the cost of a gradient-free or hybrid algorithm. Available data suggests that multimodality may be highly dependent on various parameters within the problem definition, so understanding the behaviour of the problems most often encountered in aerodynamic shape optimization is essential. For practitioners engaged in aircraft design these problems are largely transonic, but may be subject to a wide variety of constraints and permit any combination of degrees of freedom; previous studies have shed some light onto these questions, but significantly more data is required to fully understand the broader picture.

The objective of this study is to provide information to enable practitioners to make informed assessments of the risk presented by multimodality in a variety of design spaces. This is accomplished through a parametric study of multimodality in the optimization of wings subject to a variety of degrees of freedom, constraints, and operating points. A particular focus is on quantifying the likelihood that a gradient-based optimizer will converge to a local optimum that is significantly different from the global optimum. While one might argue that any finite likelihood of such an occurrence would suggest the need for a global optimizer, practical considerations may dictate that a certain amount of risk is tolerable in certain contexts.

II. Methodology

This study is undertaken using the Jetstream optimization framework, which is described in further detail below. Additionally, a discussion of the gradient-based multistart (GBMS) protocol used to examine each problem is provided, as is a description of the structure of the study and a variety of metrics used to quantify the degree and nature of the multimodality in each given problem.

A. Aerodynamic Shape Optimization framework

Jetstream utilizes Diablo, a three-dimensional structured multiblock CFD solver capable of solving both the Euler [24] and RANS [25] equations; for the RANS equations the Spalart-Allmaras one-equation turbulence model is used. Spatial discretization is accomplished using second-order summation-by-parts operators and scalar or matrix numerical dissipation. Boundary conditions and block interfaces are enforced weakly with simultaneous approximation terms. The steady-state solution is obtained with a parallel Newton-Krylov-Schur algorithm which employs an approximate-Newton phase to generate the initial iterate for the subsequent inexact-Newton phase. GMRES is used for the solution of linear systems with an approximate-Schur preconditioner.

Gradients are obtained using an implementation [26, 27] of the discrete adjoint method [1]. The adjoint systems are solved with a flexible variant [28] of the GCROT Krylov method [29]. Once calculated, the gradients are supplied to SNOpt [30], which provides the subsequent iterate using a sequential-quadratic-programming algorithm with Hessians obtained via a BFGS update method.

Mesh deformation is achieved using a linear elasticity method adapted by Hicken and Zingg [28] from the work of Truong et al. [31]. To improve efficiency, the computational mesh is fitted with a coarse B-spline control mesh, the mesh deformation being applied to this control mesh, which is then used to generate a deformed computational mesh.

B. Geometry Parameterization and Control

Geometry parameterization and control are separated in Jetstream, with parameterization occurring through B-spline patches which define the surface itself, and control provided through an axial and free-form deformation (FFD) scheme [32] depicted in Figure 1. The surface patches are fitted with 4th order B-splines; the HWB case uses 12 control points per edge while in the CRM cases each patch has between 5 and 11 control points per edge. These surface control points are embedded within B-spline volumes, referred to as FFD volumes, these volumes are deformed by manipulating their control points, which in turn deforms the underlying geometry. Cross-sectional shape changes can be realized by manipulating FFD control points, which are organized into cross-sections. Each cross section may have a twist and taper design variable, which respectively act as a local rotation and scaling factor of the entire cross section, and each cross-sectional control point may have its own section design variable to enable control of the cross-sectional shape. Large deformations are driven by a B-spline curve, referred to as the axial curve, one per FFD volume. The axial curve has its own set of control points and may be deformed by manipulating them; each cross-section is constrained to be locally perpendicular to the axial curve and so by deforming the curve, cross sections can be easily rotated or translated in complex fashions. Each axial control point may have up to three design variables - sweep, span, and dihedral - which are respectively translations of the axial control points in the chordwise, spanwise, and vertical directions.

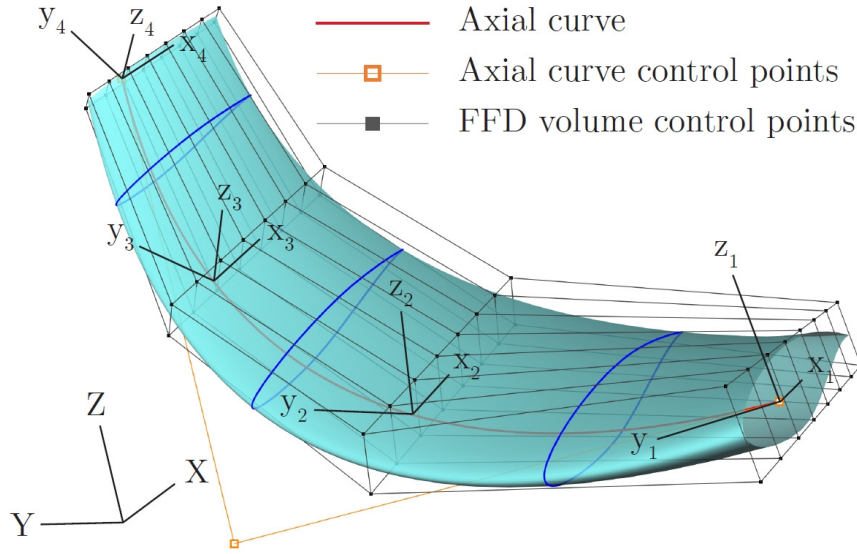


Fig. 1 Axial and Free Form Deformation Geometry Control System [32]

C. Gradient-Based Multistart Algorithm

Gradient-based multistart (GBMS) is an optimization method wherein multiple gradient-based optimizations in a given design space are run starting from different initial geometries within this space [16]. The initial geometries are generated using a sampling algorithm and together represent a sample of the design space. This method represents a compromise between the speed of gradient-based optimization and the robustness of stochastic methods. GBMS offers two key properties that make it ideal for a study such as this: first, it grants some ability to study multimodal design spaces without accepting the prohibitive cost of gradient-free aerodynamic shape optimization [16]; second, while it greatly increases the probability of locating the global optimum, GBMS also presents any other inferior local optima which are located. This provides a cost-effective means of performing a rough mapping of the design space.

The quality of the initial sample is critical to the success of a GBMS study; a robust exploration of the design space is essential to ensure that important regions are not missed. As proposed by Chernukhin and Zingg [16], sampling is accomplished using Sobol sampling [33, 34], which has several desirable qualities: it has superior space-filling characteristics to random sampling, it is deterministic for a given set of input parameters - permitting easy repetition or comparison of samples - and the position of each specific sample point is independent of the number of other sample points - allowing one to easily add additional points later without compromising the sample.

The Sobol algorithm provides a value between 0 and 1 to sample between user provided upper and lower bounds; the determination of these bounds plays an equally significant role in the final quality of the sample. In general, the region of interest for the initial geometries represents a sub-space of a larger design space; it is important that this region

be properly defined. A purely random set of initial geometries will almost certainly be composed overwhelmingly of geometries that are physically impossible, too difficult to manufacture, or do not lead to steady flow solutions. However, an unnecessarily constrained sample could exclude a viable global optimum from being located. The importance of proper sampling constraints is clear; equally important is when they are applied. The simplest method is to sample broadly and then rely on linear constraints enforced prior to the first design iteration to bring the samples into the feasible region, but this will simply snap all infeasible geometries to the feasible region boundary, dramatically oversampling the boundary while undersampling the region within it. To obtain a proper, representative sample the constraints must be applied during sampling, incorporated into the sampling constraints.

Finally, how the constraints are defined has considerable influence on the quality of the resulting sample; a naive approach would be to simply sample between the upper and lower bounds of each design variable. This comes with substantial weaknesses which are illustrated in Figure 2. Suppose the sampling is being applied to a simple wing controlled by six control points with a baseline shape as shown in Figure 2a, and suppose some large but feasible deformation like that shown in Figure 2b is desired. This geometry could be achieved by permitting large upper and lower bounds on each control point, however, as Figure 2c shows, this also permits a much larger pool of clearly unacceptable geometries, which would likely dominate the design space and almost certainly fail either in the mesh movement or flow solution stage. One can limit the number and extent of such geometries by reducing the sampling bounds on each design variable, but as illustrated in Figure 2d, this also precludes the target geometry and many other sensible geometry families. If one instead defines the reference point for the sampling bounds of each design variable as the sampled value of the previous design variable, a cascading chain of interacting constraints is created. Figure 2e shows such a system is able to permit highly flexible deformations with bounds as tight or tighter than those used in Figure 2d. This is analogous to a 3D animation of a human body permitting the foot to move anywhere in space provided it is within a certain radius of the knee joint, which is in turn constrained by a set relationship with the hip position and so on. This was the approach taken by Chernukhin & Zingg [16] in their multimodality study, and was shown to be quite effective at producing high quality, sensible sample geometries. However, the B-spline geometry control system used at the time necessitated a specialized constraint system for every class of geometry; in this work, moving to an axial FFD geometry control scheme has permitted these constraints to be generalized to an entirely geometry-independent form. The constraints need not be restricted to sampling, the idea of converting optimization bounds from an absolute reference to a relative one is an interesting area of future exploration.

D. Study Structure

This study seeks to obtain one critical piece of data from each examined design space: how much risk is associated with neglecting multimodality in this design space, if any. Careful design of experiments is necessary as there is no definitive way to determine when one has sufficiently explored the design space. In an effort to balance thoroughness

and cost a two-step structure is adopted. For each case a large sample of initial geometries is generated, then a set of 17 initial geometries - 16 sample points and the baseline - is optimized to completion. The resulting geometries are then examined for any indication of multimodality - the criteria used are provided in Section II.E. If multimodality is located, the case is terminated, otherwise, a second set of 16 samples is attempted; if no additional optima are located in this set the problem is concluded to be unimodal for the purposes of this study. Because not all initial geometries successfully converge, the final number of converged geometries is often somewhat less than 17 or 33; in cases where convergence proves particularly difficult, additional samples are attempted until at least 10 converged geometries were located. Thirty-three initial geometries are not enough to constitute an exhaustive search of a design space, particularly for the highly-dimensional, exploratory problems examined later in this study. Therefore, the objective of this study is explicitly not to locate the global optimum, but to ascertain the presence of multimodality and, if present, some sense of its degree. With these limited sample sizes, we cannot claim to have found the global optimum with any certainty, and the possible presence of additional local optima with superior performance to the best performing local optimum found cannot be ruled out.

E. Multimodality Metrics

As a result of the limitations of finite optimization convergence, in some cases the determination of whether two geometries represent distinct local optima can be difficult. Ascertaining at which point differences are sufficiently large to indicate distinct optima in a systemic manner requires a clear set of criteria that can be applied consistently. For the purposes of this study, two final geometries are considered distinct local optima if both have converged - defined as a feasibility of 10^{-5} or less, one to two orders of magnitude reduction in optimality* - and a merit function which is changing by less than 1×10^{-5} , roughly 0.03 drag counts - and satisfy at least one of three additional requirements:

- 1) final drag coefficients differing by 1 drag count or more, or
- 2) a root mean square (RMS) difference in surface shape of at least 0.05 units, or
- 3) an RMS difference in twist distribution of at least 0.2 degrees.

The 1 drag count threshold was introduced to avoid identifying incompletely converged geometries as local optima. With this threshold any local optimum that is missed does not produce significant performance variation. Our conclusions with respect to the risk of multimodality are not sensitive to the precise choice of this threshold. The third criterion ensures that distinct twist distributions are captured even if there are very small differences in surface shape and drag. A number of metrics are employed to quantify any multimodality found with the above criteria. The first three are straightforward measurements: the number of local optima, the best performance located, and the difference between the best and worst performance - the latter two in drag counts. The fourth metric, R , is a measurement of the risk associated with multimodality in a given design space and is defined as the percent of initial geometries that converged to local

*In optimization with the RANS equations, it is often not possible to achieve deeper convergence

Table 1 Classification of design space multimodality

R	Classification
$R = 0\%$	No noted multimodality
$0\% < R \leq 5\%$	Slightly multimodal
$5\% < R \leq 20\%$	Somewhat multimodal
$20\% < R \leq 50\%$	Moderately multimodal
$R > 50\%$	Highly multimodal

optima that are outperformed by the best local optimum by at least one drag count. It thus quantifies the likelihood that a local optimizer starting from a single initial geometry will get trapped in a local optimum significantly inferior to a best estimate of the global optimum.

Using R , the nomenclature of Chernukhin & Zingg [16] for categorizing degrees of multimodality is adapted for the test procedure used in this study, categorizing the degree of multimodality based on the risk rather than the raw number of local optima. The thresholds used in this work are given in Table 1; unlike Chernukhin & Zingg’s work, there is no explicitly unimodal category. This is because given the finite sample sizes involved it is not possible to conclude definitively in this study that a problem is unimodal, only to recognize a shrinking likelihood of locating additional optima and the weak nature of multimodality in that search space, if any.

The value of R is not sufficient to quantify the shape of a design space, nor the risk posed by multimodality within it. For instance, the distribution of local optima within the design space - how much they underperform the best optimum and how many initial geometries converge to them - impacts the overall risk of overlooked performance associated with the use of a single initial condition. Additionally, researchers engaged in aerodynamic shape optimization, particularly exploratory optimization, are not necessarily singularly focused on maximizing performance. It is often desirable to determine the geometric variation across the local optima in the design space, in which case geometrically distinct, difficult to locate local optima are equally or more problematic than easy to find, poorly performing optima. Previous work [35] has attempted to quantify these qualities, but for simplicity here the focus is placed on R as a measure of risk, and further insight can be obtained by looking at the range metric in conjunction with R . We are measuring the risk in terms of the likelihood that a reasonable arbitrary initial geometry will converge to an inferior local optimum. Another way of defining the risk is in terms of the likelihood that the baseline geometry will converge to an inferior local optimum, as opposed to an arbitrary initial geometry. In an aerodynamic shape optimization problem typical of detailed design where the final geometry is likely to be quite similar to the baseline geometry, the latter risk is substantially lower than the former. However, in preliminary or exploratory design, the final geometry is expected to be substantially different from the baseline geometry, and thus the risk associated with starting with the baseline geometry is no different from that associated with starting with any reasonable arbitrary geometry.

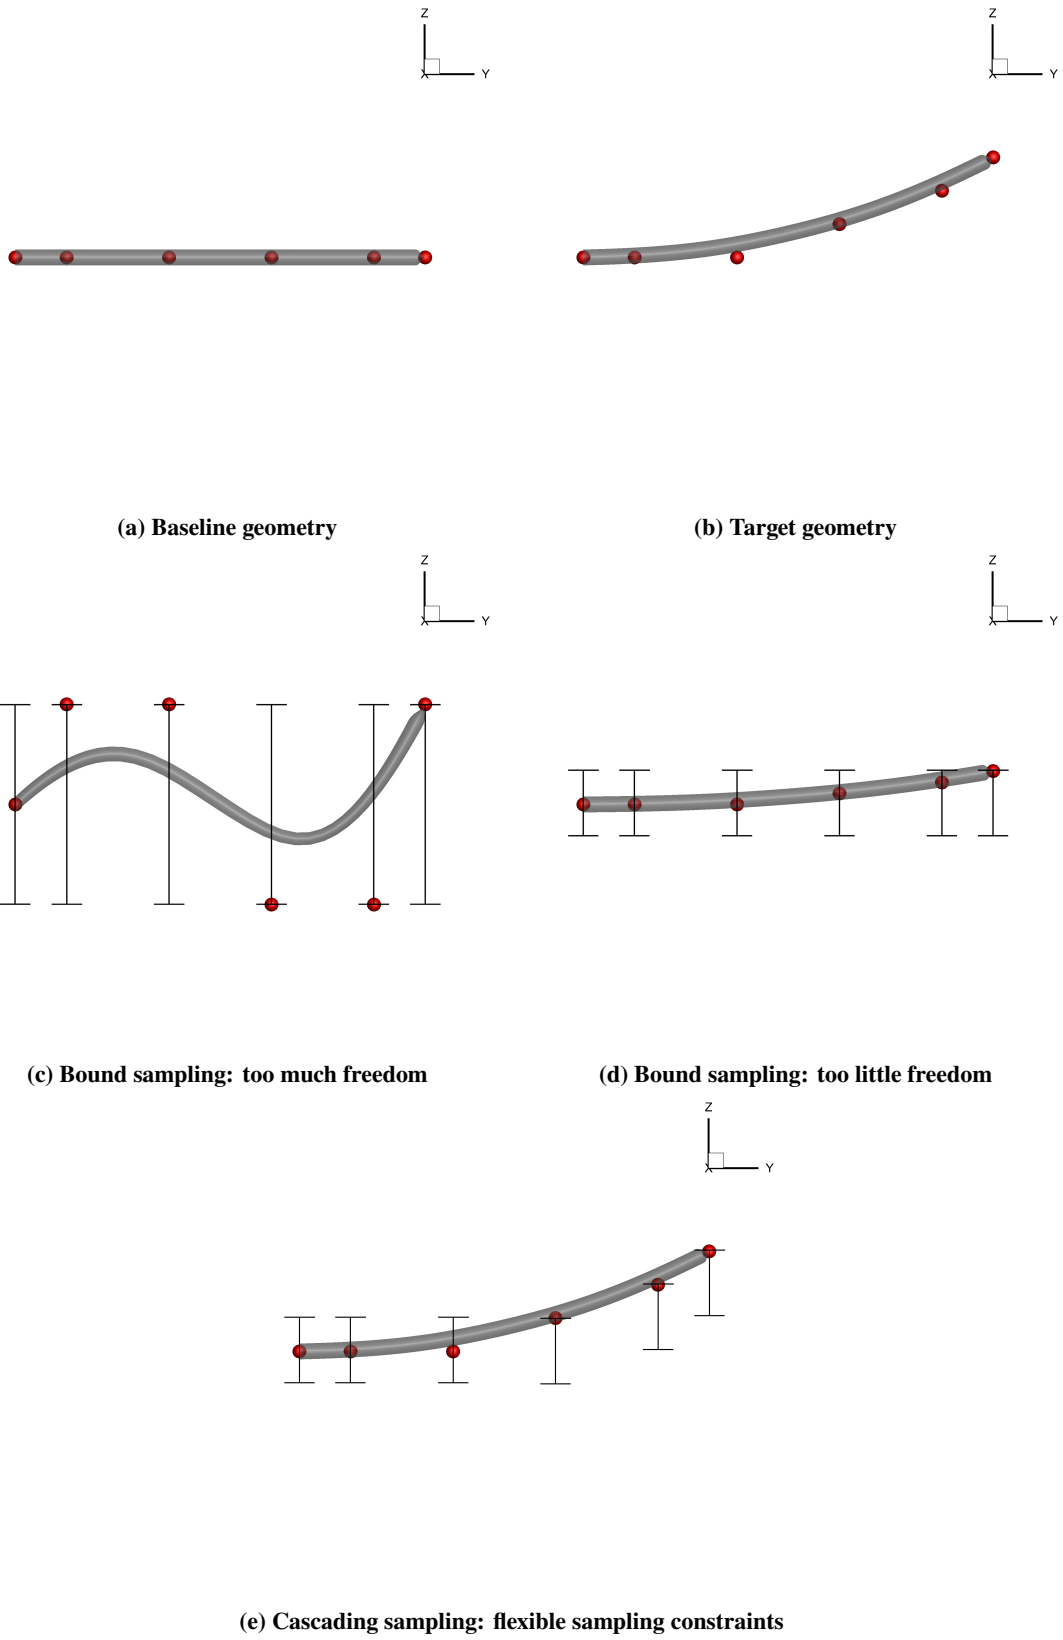


Fig. 2 Various constraint methods for sampling. Control points in red, bounds in black.

III. Results

The effect on multimodality of optimization degrees of freedom (DOF), constraints, numerical mesh, and operating points is examined. A note on nomenclature is relevant here: in this work there must be a differentiation between the number of design variables - the number of values the optimizer may use to improve the design - and a more general indication of the type of design variables that are permitted, such as sweep or twist. To aid in this distinction slightly non-standard terminology is employed. From this point “design variable” will be taken in the traditional sense within optimization, while “degrees of freedom” will be used exclusively to refer to types of design variable. For example, permitting twist to vary adds a single DOF, but potentially many design variables.

Five distinct studies are presented, each with multiple variations, ranging from practical design optimization to highly exploratory cases. All cases are governed by the RANS equations. The first case examines two variations of ADODG case 5, the twist and section optimization of the CRM wing-only geometry. In the second and third studies this baseline geometry is retained, but the number of degrees of freedom and constraints are varied to determine their impact on multimodality. The fourth study examines the effect of multiple operating points on two variations of this case, while the fifth study exchanges the CRM problem for an HWB optimization based on the work of Reist et al. [3] in order to examine multimodality in the context of a problem more representative of exploratory optimization. Finally, a brief investigation of grid dependence in multimodality is undertaken.

A. CRM ADODG Case

This case is defined as the Aerodynamic Design and Optimization Discussion Group (ADODG) case 5 and represents a common family of wing optimization cases, performed in viscous flow at a Mach number of 0.85 and a Reynolds number of 5 million. The Common Research Model (CRM) wing-only geometry, shown in Figure 3, is used as the baseline geometry with dimensions scaled by the mean aerodynamic chord (MAC). The computational mesh is a modified version of that used by Koo and Zingg [18] with parameters as listed in Table 2. This grid has been shown to be sufficiently accurate for aerodynamic shape optimization, and a grid refinement study was not repeated here as the specific drag values are not of primary interest in this study. Geometry control is accomplished with the FFD system illustrated in Figure 3c. Two separate FFD volumes are used, meeting at the crank, the inboard with three spanwise stations and the outboard with five. Each FFD volume has its own axial curve, fitted to the wing trailing edge, which can provide global control. In this case only twist and section design variables, in addition to angle of attack, are available to the optimizer while all other design variables are fixed, totalling 147 active design variables. The optimization problem itself can be formally stated as

Table 2 CRM mesh parameters

Nodes	1,068,856
Blocks	40
Off-Wall Spacing	7.3×10^{-7} MAC
y^+	0.13

$$\begin{aligned}
& \text{minimize} && C_D \\
& \text{w.r.t.} && v \\
& \text{subject to} && C_L = 0.5 \\
& && V \geq 0.2617 \text{ MAC}^3 \\
& && C_M \geq -0.17
\end{aligned} \tag{1}$$

where C_D is the drag coefficient, C_L is the lift coefficient, V the wing volume, and C_M the pitching moment coefficient. Apart from generous upper and lower bounds on each design variable, no linear constraints are enforced. This test was examined using the sampling algorithm and test protocol outlined previously. Ultimately 33 initial geometries were attempted, 29 of which converged, all to a single optimum. The full results are provided in Table 3 and uniformly depict a problem for which multimodality presents a negligible risk, in agreement with previous studies of this problem [17, 18].

The problem was repeated with the addition of a minimum thickness constraint - denoted as CRM(MT) - requiring the thickness at each section to be not less than 85% of the baseline thickness at that section. With the addition of the minimum thickness constraint, once the 15 successful initial geometries had converged a second local optimum was noted, so the test did not proceed to a full 33 initial geometries. Based on the R value of 7% and performance range of 7.6 drag counts this problem is deemed to somewhat multimodal with a small but nonnegligible risk of significant lost performance if multimodality is ignored.

Figure 4 plots the twist distributions of the identified local optima from both cases relative to the freestream flow direction. The best twist distributions in both cases, shown in black, are highly similar. The addition of the minimum thickness constraint has not impacted this optimum, but introduced a new, inferior but locally optimal twist distribution that is not observed without the minimum thickness constraint. An alternative explanation is that the supposed new optimum was in the original design space, but was simply missed. To investigate this, the second optimum from the thickness constrained case was re-optimized from its final position with the minimum thickness constraint disabled. Doing so caused the geometry to converge rapidly to the optimum found without the minimum thickness constraint. This is evidence that a local optimum similar to that found with the thickness constraint present does not exist in its

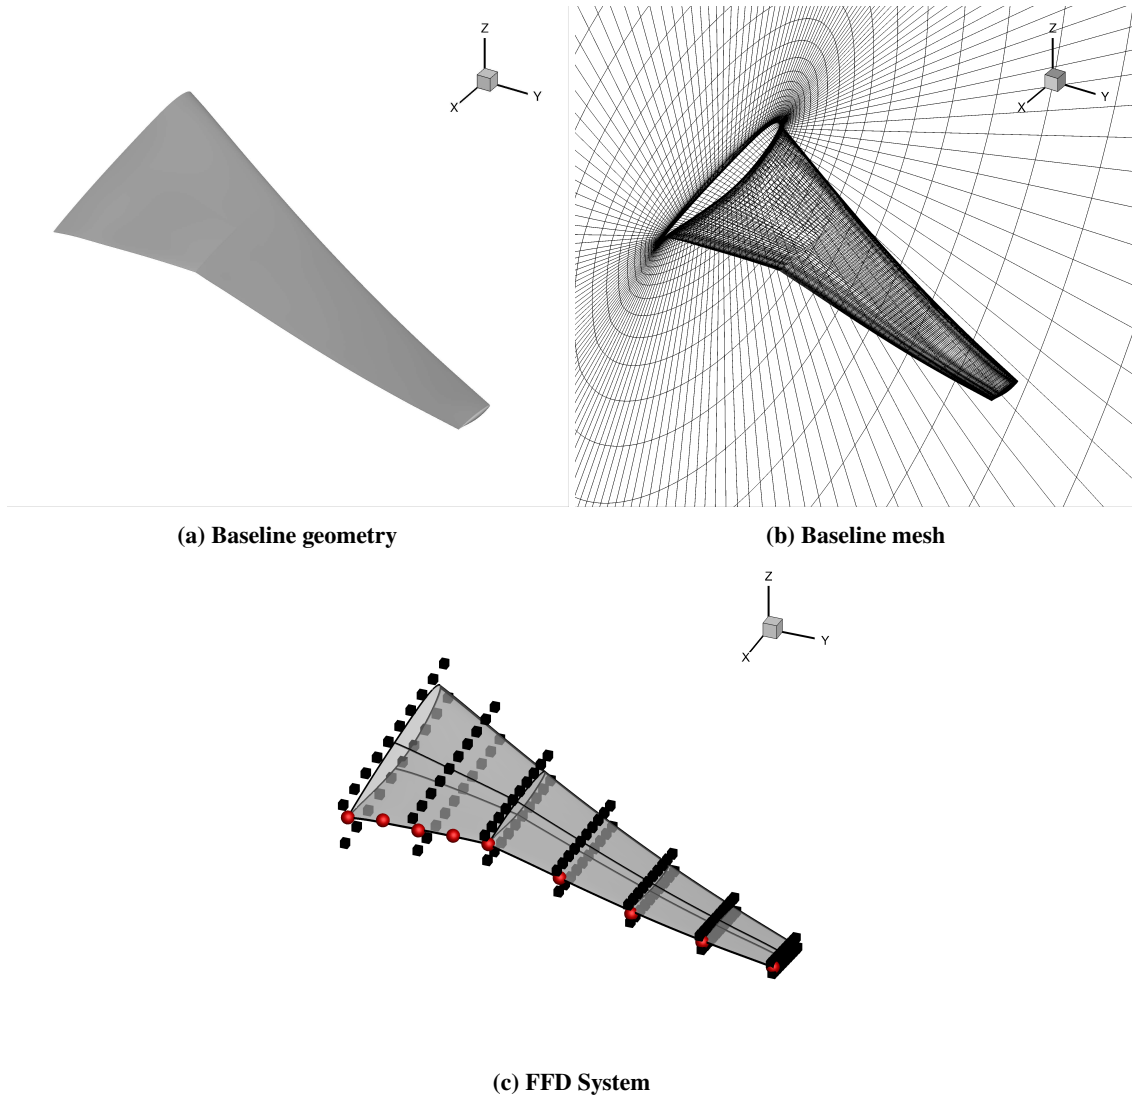


Fig. 3 Geometry, mesh, and control for CRM baseline wing


absence, and that this minimum is associated with the minimum thickness constraint.

B. DOF Study Under Transonic Conditions

One of the questions of most immediate interest regarding multimodality is how it scales with DOF, a question of particular importance given the growing study of high-dimensional, exploratory optimization problems. Using the baseline geometry in Figure 3 and the problem definition in Equation 2 the available degrees of freedom are gradually increased, starting with solely section shape and sweeping up to a highly dimensional case permitting large scale deformations of the geometry. The cases were selected to be representative of both preliminary and detailed design. Cases more typical of conceptual design were not a focus as these are not generally examined with high-fidelity tools. Each case is denoted by a five digit code of the form xxx_xx , where each digit refers to a potential degree of freedom

Table 3 ADODG CRM multimodality study results

Test	Successful Tests	Optima	Best (counts)	Range (counts)	R
CRM	29	1	191.9	0.0	0%
CRM(MT)	15	2	194.5	7.6	7%

**Fig. 4 Spanwise twist distributions for ADODG CRM wing-only cases. Best optimum in black**

and takes either a 0 or 1 to denote whether that freedom is active. From left to right each digit corresponds to twist, taper, section, and then the two examined axial degrees of freedom - sweep and dihedral. A quick reference guide for which DOF are active in each case is provided in Table 11 in the Appendix. No additional constraints are enforced, apart from a projected area constraint requiring $S = 3.407$ where applicable. It should be noted that the 101_00 case corresponds to the standard ADODG case examined above. The complete results are provided in Table 4.

Based on this data, the risk of multimodality scales with DOF. The first five cases in Table 4 permit only various combinations of section shape, twist, and chord modification and present with a low to moderate risk of lost performance from multimodality. An appreciable range of performance values are apparent even in these relatively limited problems, with ranges as low as 0.1 drag counts but as high as 3.5 in one case - however, the low risk value in the latter case indicates that this outlying local optimum is attracted by a small region of the design space, and so is less likely to accidentally trap a local optimizer. Risk too shows a wide range of values across these cases; one presents without any detectable multimodality, another is slightly multimodal, while the remaining three cross-sectional cases are somewhat multimodal, two of those approaching the upper end of the “somewhat multimodal” range in Table 1. These results show that there can be an appreciable risk of lost performance from multimodality even in cases with limited geometric freedom. The risk increases when the wing shape is permitted more geometric freedom, as shown by the last two cases

Table 4 Degree of Freedom multimodality study 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. DVs are the total number of design variables.

Test	DVs	Successful Tests	Optima	Best (counts)	Range (counts)	R
001_00	141	12	2	195.1	1.3	17%
011_00	148	13	2	181.8	3.5	8%
101_00	147	29	1	191.9	0.0	0%
110_00	14	33	2	199.7	0.1	3%
111_00	154	11	3	181.0	1.0	18%
111_01	162	14	4	176.2	4.6	43%
111_11	170	10	7	174.1	24.6	90%

in Table 4, which permit all of the previously discussed DOF but add, respectively, dihedral and sweep with dihedral. In both of these cases a steep growth in risk is observed. Relative to 111_00, adding dihedral control as in 111_01 produces a nearly five-fold increase in performance range and more than double the R value, a moderately multimodal design space with large potential for lost performance. Adding sweep in addition to dihedral in 111_11 increases both metrics again by similar proportions and yields a highly multimodal design space. This increased multimodality in preliminary design studies can be observed in Figures 6, 7, and 8; it should again be emphasized that the lack of variable planform shapes in 111_00 and 111_01 does not conclusively prove the non-existence of such optima, but does suggest that if they exist they are not highly dominant in the design space.

The importance of sufficient optimization convergence in a study such as this cannot be overstated, as insufficient convergence will artificially inflate the apparent multimodality in the design space with false local optima. An exhaustive presentation of the convergence histories of all cases is not practical here, but for illustrative purposes convergence histories for representative cases are provided in Figures 5a to 5g. These highlight the convergence which was required in all cases, characterized by an optimality reduction of one to two orders of magnitude, feasibility of 1×10^{-5} or less and the merit function changing by fewer than 0.03 drag counts.

The first five cases in Table 4 are typical of aerodynamic shape optimization problems arising during the detailed design phase, having limited geometric freedom. For all of these cases, the optimization converges to the best local optimum from the baseline CRM geometry, which is a well designed aircraft wing with limited room for improvement. This supports the idea that the risk associated with multimodality is low for such problems, although it is important to bear in mind that the best local optimum found via the limited sample size used here is not guaranteed to be the global optimum. It is expected that the behaviour of the baseline geometry is not specific to the CRM, but to any case where a well-designed initial geometry is available and is likely to lie close to the best optimum in the design space. If this is true then as design spaces become larger and expected variations from the initial design increase, then the ability of

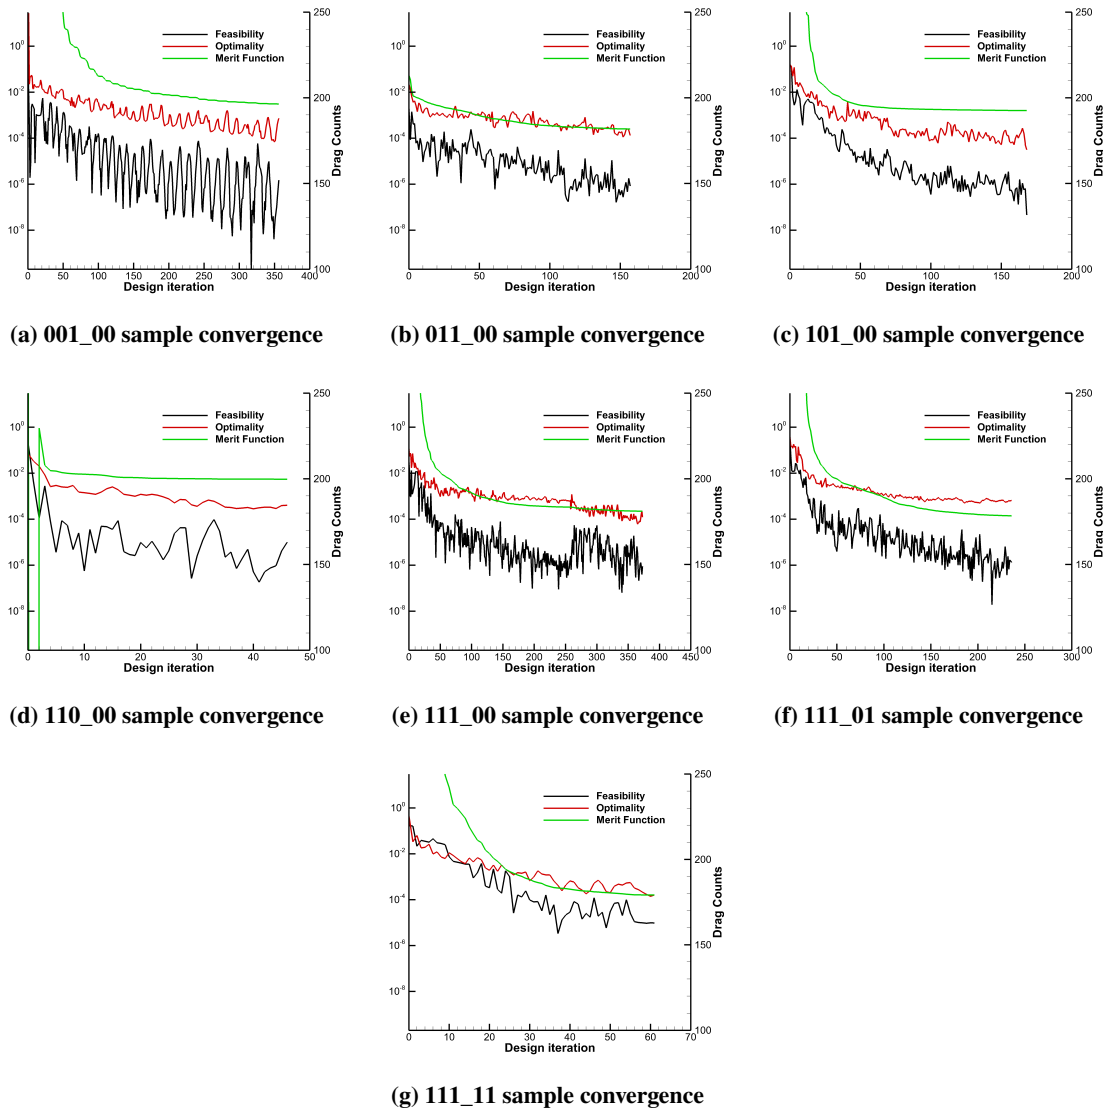


Fig. 5 Representative convergence histories

even a relatively well-engineered design like the CRM to consistently locate the best local optimum would be expected to decrease. This is exactly what is observed in last two cases in Table 4, which are more typical of preliminary design and for only one of which the baseline geometry locates the best local optimum. While this is only a sample size of two, it will be shown in the following section that the influence of the baseline geometry is consistently reduced in preliminary design problems.

Therefore in detailed design problems primarily modifying cross-sectional shape or twist there is still the potential for lost performance due to multimodality; however, this risk is reduced by the common availability of a well-designed initial geometry. When larger shape changes are permitted, failing to adequately address multimodality significantly reduces the probability of locating the global optimum. Beyond these general trends one must be very cautious extracting

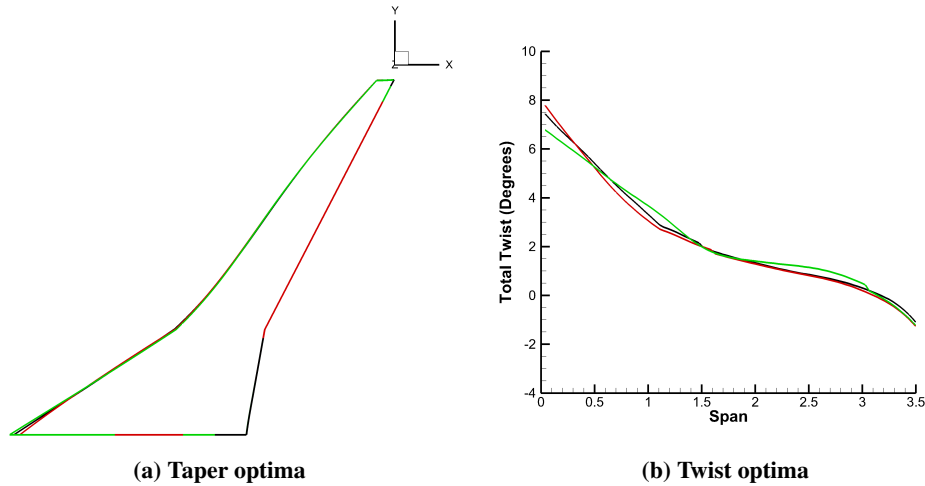


Fig. 6 Local optima for 111_00 DOF case. Best optimum in black.

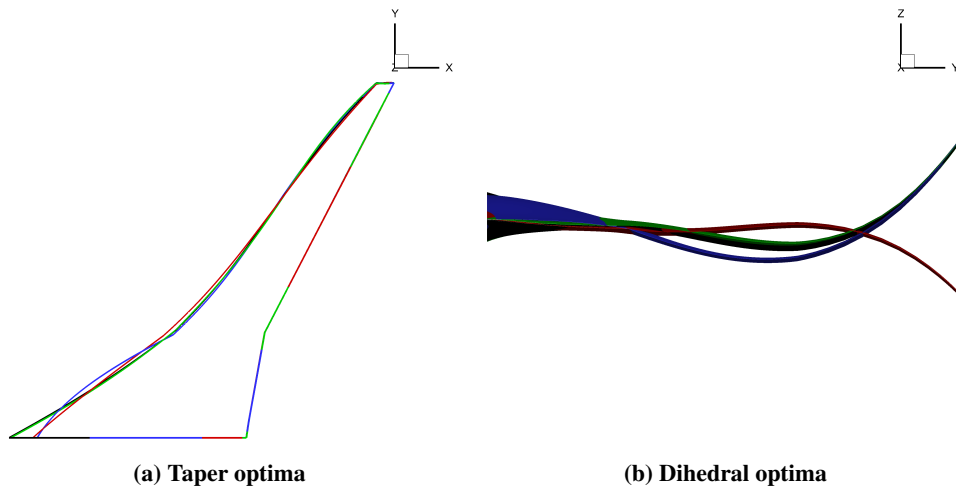


Fig. 7 Local optima for 111_01 DOF case. Best optimum in black.

specific relationships from the data. There are no clear trends between specific DOF and the amount of risk in a design space, making an a priori assessment of the risk of multimodality in a given optimization problem difficult. Nevertheless, the general trends observed here give the designer a good idea of what to expect with respect to the risk of multimodality in the optimization of wings.

C. Constraint Study

The previous study has shown significant potential for multimodality in high-dimensional, exploratory aerodynamic shape optimization. However, while exploratory optimization cases are characterized by large geometric freedom, they are also often subject to a number of constraints which could potentially have a major impact on the degree of multimodality found in the design space. For this reason, a detailed study of constraints in transonic optimization is presented.

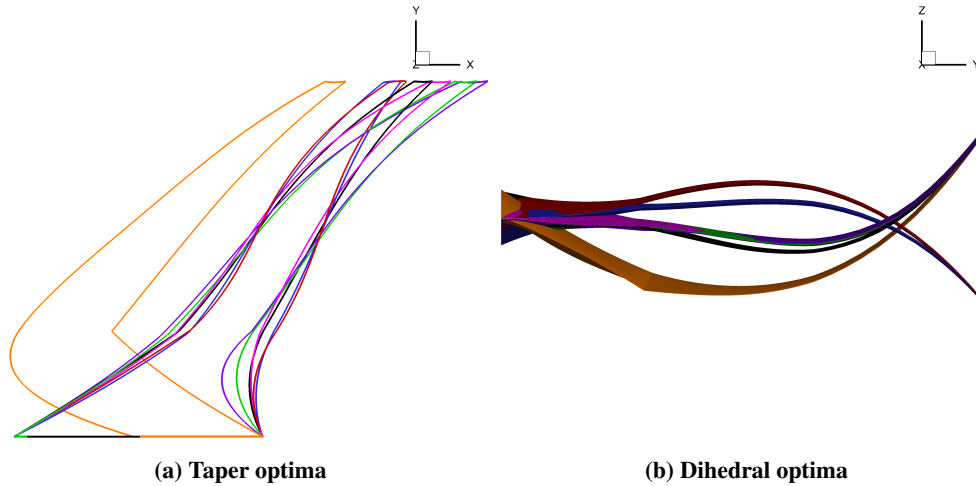


Fig. 8 Local optima for 111_11 DOF case. Best optimum in black.

This study uses the 111_11 case, with 170 design variables, from the DOF study as a baseline and varies the constraints enforced. Each case is again denoted by a five digit code of the form xxx_xx where each digit corresponds to a degree of freedom - twist, taper, section, sweep, and dihedral - though here each digit can take one of several values to indicate what, if any, constraint is enforced on that degree of freedom. “L” defines a linear distribution, “C” requires a constant distribution, and “M” enforces a strict minimum value. A “P” appended to a case denotes that the pitching moment constraint is not enforced. The additional constraints active for each case are listed in Table 12 in the Appendix. It should be noted that in the case of linear distributions, they are required to be linear only within a given FFD structure. As the two constituent FFD volumes meet at the crank of the baseline geometry this permits discrete inboard and outboard distributions. The tests are further broken down into four broad categories based on the degree of constraint placed on the leading and trailing edges of the wing, these are: baseline cases, with few if any constraints whatsoever, cases with no constraints on the leading and trailing edges, cases with some edge constraints - requiring them to be straight in either the dihedral or sweep dimensions - and cases with “full” edge constraints - requiring them to be entirely straight, save for the crank. The leading and trailing edge constraints are highlighted for two reasons: first, requiring a straight leading or trailing edge is a common quasi-structural constraint during practical optimization; second, since these constraints represent a significant limitation on two of the clearest contributors to multimodality - sweep and dihedral - it is useful as a proxy for the overall “amount” of constraint in the problem. The results are broken down along these lines in Table 5, along with averages for each category of problem.

Table 5 Constraint study results: $M=0.85$, $Re=5$ million. 000_00 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 ‘-P’ denotes that a pitching moment constraint is not enforced.

Test	Successful Tests	Optima	Best (counts)	Range (counts)	R
Baseline Cases					
000_00	10	7	174.1	24.6	90%
000_00-P	14	12	173.7	156.3	93%
Average	12	10	173.9	90.5	92%
No LE/TE Restrictions					
L00_00	11	7	174.4	6.8	91%
0L0_00	12	8	174.1	39.2	83%
00M_00	12	6	178.6	18.7	67%
0LM_00	15	11	180.3	18.7	87%
LL0_00	11	7	174.4	13.0	91%
LLM_00	12	8	180.3	42.6	67%
Average	12	8	177.0	23.2	81%
Some LE/TE Restrictions					
000_C0	12	5	175.9	19.4	42%
000_0C	11	6	178.4	94.5	45%
0L0_0C	14	7	179.6	23.3	93%
0L0_C0	10	6	176.9	13.9	80%
0LM_C0	10	7	190.2	48.0	60%
LL0_0C	11	5	179.4	30.9	36%
LL0_C0	11	6	177.2	24.0	91%
Average	11	6	179.7	36.3	64%
Straight LE/TE					
0L0_CC	15	6	181.0	26.6	67%
LL0_CC	15	5	180.9	29.2	47%
LLM_CC	16	2	195.6	17.7	50%
Average	15	4	185.8	24.5	55%

This study was in part motivated by the interesting result observed earlier, where multimodality was introduced into the design space by enforcing a stricter minimum thickness constraint on the ADODG case 5 problem. Examining the average values reported in Table 5 one can see that, by certain metrics, the opposite trend is apparent here. Moving from the baseline cases through the Straight Leading Edge/Trailing Edge (LE/TE) results at the bottom of the table, there are steady reductions in the number of optima and in R , together suggesting a declining risk from multimodality as large planform shape changes become increasingly restricted. However, no such trend is visible with the performance range, which varies randomly across the reported averages. The lack of a decline in range along with the reduced number of optima and smaller R values suggests that the apparent decline in risk may be deceptive, driven largely by relatively clustered optima coalescing into a single point while significant outlying optima remain largely unaffected. This behaviour can be observed visually by plotting the local optima for the three Straight LE/TE cases, as in Figures 9 and 10. While the number of optima varies across the three sub-problems under discussion, the geometric families represented remain consistent: forward- and back-swept wings with clear breaks at the crank and dihedral distributions ranging from flat to clearly negative. The flow direction in these figures is from left to right, so observant readers will notice a consistent tendency in the straight LE/TE cases to remove the trailing edge crank in favour of a less-common leading-edge crank. This is a side-effect of the constant sweep constraint, which is enforced at the trailing edge and therefore forces the trailing edge to be straight, while the optimizer can insert a leading edge crank as a result of the linear taper constraint, which permits different taper ratios in the two adjoining FFD volumes. While such designs may not be desirable in backswept configurations, this has no effect on the multimodality in these cases nor on our conclusions drawn regarding them. The geometric consistency seen in these cases reinforces that the constraints are primarily reducing multimodality by consolidating nearby optima without fundamentally changing the shape of the design space. The persistent presence of multimodality in optimization problems with large degrees of geometric freedom, even when substantially constrained, has significant implications for exploratory optimization. It is likely, based on the available data, that many such problems will be unavoidably multimodal.

The LL0_CC case, despite being contained within the 0L0_CC design space, locates a slightly better local optimum than the latter. This is merely due to the finite nature of the number of initial geometries used, which cannot guarantee that every optimum within the design space will be located; starting an 0L0_CC optimization with the best LL0_CC optimum as the initial geometry rapidly converges to a superior local optimum that was not found by the original set of 15 initial geometries. This should serve to reinforce the nature of these studies as representative, rather than exhaustive. The data presented here, and in all the cases in this work, is therefore best viewed in the context of the discussed trends: even extremely strict constraints are ineffective at ameliorating the risk associated with multimodality in high dimensional design spaces. The one potential exception is the pitching moment constraint, as its removal generates a marked increase in multimodality across every reported metric. This suggests that flow-dependent constraints have a greater impact on multimodality than geometric constraints, but a more thorough exploration of this question must await




Fig. 9 Planform shapes for straight LE/TE cases. Best optimum in black.

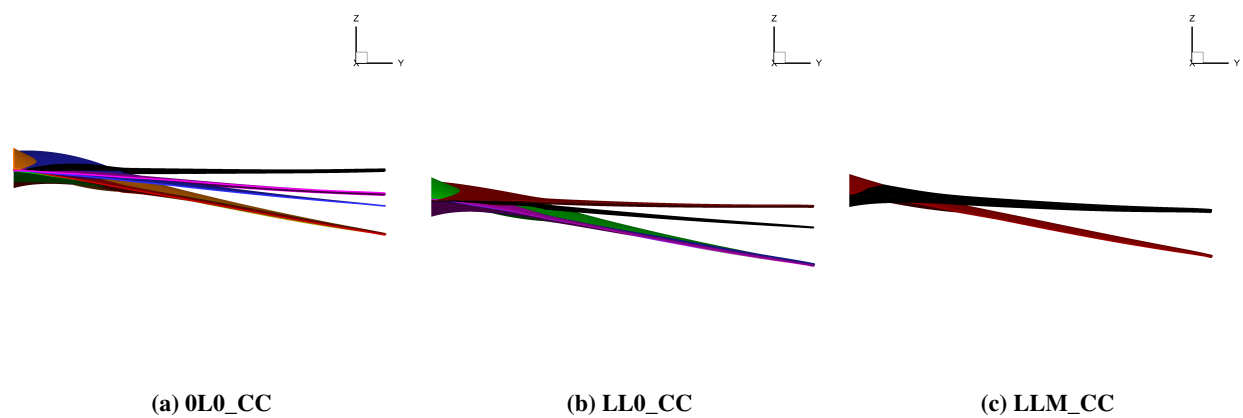


Fig. 10 Dihedral shapes for straight LE/TE cases. Best optimum in black.

a future work.

Returning briefly to the discussion in Section III.B regarding the behaviour of the baseline geometry, these results further reinforce the conclusions drawn there. In this sample of optimization problems that are typical of preliminary design in terms of their geometric freedom, the ability of the CRM initial geometry to locate the best optimum reliably is significantly reduced; its failure rate has grown from one in seven in the DOF study to six out of eighteen here. For such problems, the risk incurred by performing a single optimization from one initial geometry, even a well-designed one, is substantial and can be greatly reduced through a multistart approach.

Table 6 Multipoint CRM results: three operating points

Test	DVs	Successful Tests	Optima	Best (counts)	Range (counts)	<i>R</i>
CRM (MT) Case						
Single-Point	147	15	2	194.5	7.6	7%
Multi-point	147	26	1	202.9	0.0	0%
LLM_CC Constraint Case						
Single-Point	170	16	2	195.6	17.7	50%
Multi-point	170	11	3	202.8	25.6	45%

D. Multi-point CRM

One largely unresolved region of the literature is the impact of multi-point optimization on multimodality. To study this, two cases - the CRM (MT) case from the ADODG Case 5 study and the LLM_CC constrained case - were re-optimized as multi-point problems. Each used the same three operating points: $M = 0.82$, $M = 0.85$, $M = 0.88$. The Reynolds numbers at each point are 4.82 million, 5 million, and 5.18 million, with lift targets of 0.537, 0.500, and 0.466. Additionally, the pitching moment constraint is applied at the second operating condition. These operating points were previously applied to the CRM geometry by Lee et al.[36]. Results from both studies are provided in Table 6; as well the results from single point versions of both cases are replicated. In the multi-point cases performance as reported for the Best and Range metrics, as well as that used in R are calculated using the average drag across all three operating points.

The provided results show that while multi-point optimization may change the nature of any multimodality present in the design space, it does not appear to play a significant role in either expanding or mitigating the risk presented by that multimodality. The single-point CRM (MT) case was on the low-end of the “somewhat multimodal” range while the multi-point case presents no apparent multimodality; however, this reduction was entirely due to the elimination of a single, rare local optimum and both design spaces are heavily dominated by the best optimum. In the LLM_CC case, the multi-point problem introduced a new local optimum, increasing the number from two to three and the range by nearly eight drag counts. However, the risk declined slightly, from 50% to 45% - in the single-point case the best optimum was found by eight out of 16 initial geometries, but six out of 11 in the multi-point case. Overall, from a purely risk-based perspective, the additional local optimum does not render the design space much more difficult to navigate than it already was for a gradient-based optimizer, though it should be noted that the increased performance range in the multi-point case increases the maximum potential lost performance represented by R , if not the value of R itself.

Therefore, these results lead to the conclusion that multipoint optimization does not have a large impact on multimodality-associated risk in gradient based optimization, though it may impact gradient-based optimization in other ways by changing the shape of the design space.

Table 7 HWB mesh parameters


Nodes	1,220,736
Blocks	128
Off-Wall Spacing	4.9×10^{-7}
y^+	1.2

E. Hybrid Wing Body

This case is based on work by Reist et al. [3] and explores the lift-constrained drag minimization of the baseline Hybrid Wing Body (HWB) shown in Figure 11, subject to moment, volume, and cabin shape constraints. The numerical mesh is shown in Figure 11b with parameters as tabulated in Table 7. All dimensions are scaled by the MAC and optimization occurs at a Mach number of 0.78 with a Reynolds number of 76 million. The optimization problem may be formalized as


$$\begin{aligned}
& \text{minimize} && C_D S \\
& \text{w.r.t.} && v \\
& \text{subject to} && C_L S = 0.12 \text{ MAC}^2 \\
& && V \geq 0.0786 \text{ MAC}^3 \\
& && C_M S = 0
\end{aligned} \tag{2}$$

where S is the projected area. A cabin shape constraint requires that the shape of the body region of the aircraft be able to contain a cabin of pre-defined shape and size. This ensures that there is sufficient cabin volume to satisfy passenger capacity requirements. The geometry control system is shown in Figure 11c; as axial freedom is not used in this case the axial control points are not displayed. Twist, taper, and section shape modification are permitted throughout the geometry, while sweep, span, and dihedral are held fixed, yielding 350 design variables. A linear taper distribution is required everywhere except for the transition region between the fuselage and inner wing. The results from this study are reported in Table 8 : four local optima were located within the first 16 converged initial geometries and so the sample was not expanded to a full 33 points. A full quarter of the converged initial geometries located optima significantly underperforming the minimum, yielding an R value of 25% and a moderately multimodal problem. Combined with a performance range of over four drag counts, this is a problem for which ignoring multimodality carries considerable risk.




(a) Baseline geometry

(b) Baseline mesh



(c) FFD system

Fig. 11 Geometry, mesh and control for HWB baseline design



(a) Twist distributions

(b) Planform shapes

Fig. 12 Local optima for HWB case. Best shown in black.

Table 8 Hybrid Wing/Body results: $M=0.78$, $Re=76$ million.

Successful Tests	16
Optima	4
Best (drag counts)	96.6
Range (drag counts)	4.4
R	25%

Figure 12 shows that multimodality in taper is largely confined to the region connecting the fuselage and the wing where the linear taper constraint is not applied, while large amounts of multimodality are found in twist distributions. Multimodality had previously been noted in more exploratory, inviscid variations of this problem [16], and the degree of multimodality seen here is compatible with that seen in the comparable 111_00 wing optimization case from the earlier DOF study. This is further evidence that multimodality in aerodynamic shape optimization is not a quirk of a particular problem, nor a particular mesh, but is endemic in many optimization problems, and as demonstrated earlier, only becomes more prominent as more freedom is allowed.

IV. Mesh Sensitivity Study

It is known that optimization in general can be sensitive to the computational mesh used, particularly if the mesh does not accurately capture the physics of the problem. This study investigates whether similar sensitivity exists for multimodality in aerodynamic shape optimization. To do so, GBMS studies were repeated on two different cases with the computational meshes refined by a factor of 1.26 in each direction, producing a finer mesh with parameters as listed in Table 9. The two cases are the ADODG(MT) case from the ADODG study and the LLM_CC case from the constraint study; these represent very different problems, with different DOF and constraints, and while both present as multimodal the nature of that multimodality is extremely different. The former case found a single additional optimum, located by just a single initial condition, while the latter was evenly split between two significantly different local optima with highly differentiated geometries.

The top-level results for both fine-mesh cases and their coarse-mesh equivalents - re-solved on the finer mesh - are given in Table 10. In both test cases there is excellent agreement across both meshes on the geometry and performance of the best optimum, number of located optima, and R , and for the LLM_CC case the secondary local optimum is nearly identical as well. One outlier in the data is the secondary local optimum from the ADODG(MT) case. While on both meshes one outlying optimum is located by a single initial geometry, the performance of that optimum is considerably worse when optimized on the coarse mesh than the fine one. It is apparent that the multimodality in these cases is not an

Table 9 CRM fine mesh parameters

Nodes	1,983,520
Blocks	40
Off-Wall Spacing	6.6×10^{-7} MAC
y^+	0.09

Table 10 Optimization mesh independence results. All values computed on fine mesh.

Test	Successful Tests	Optima	Best (counts)	Range (counts)	R
ADODG(MT)					
Fine	16	2	192.7	3.5	6%
Coarse	15	2	192.7	7.3	7%
LLM_CC					
Fine	15	2	193.0	17.7	53%
Coarse	16	2	193.4	18.1	50%

artifact of the mesh and its overall characteristics are not highly sensitive. Nevertheless, it is possible for mesh density to impact the shape of the design space, particularly in outlying regions like that in the ADODG(MT) case, and even at refinement levels more than sufficient for effective optimization.

V. Conclusions

Multimodality is often present in aerodynamic shape optimization problems typical of detailed, preliminary, and exploratory design. In problems typical of detailed design, the risk presented by this multimodality lies between 0 and 17%, typical of slightly to somewhat multimodal problems, though this is mitigated by the consistent ability of a well-designed initial geometry, which is often available for such problems, to locate a high-quality local optimum. In problems typical of preliminary design, which are characterized by increased geometric freedom and consequently a greater difference between the baseline geometry and the best local optimum, the risk associated with multimodality is much higher, ranging from 18-90%, which in our classification system identifies most of these problems as moderately to highly multimodal. Therefore, in the case of exploratory design of novel unconventional aircraft, where design experience is limited and a great deal of geometric freedom must be permitted, multimodality is very likely to occur. For preliminary and exploratory design, allocating resources to an examination of multimodality can significantly reduce the risk of overlooking a superior optimum. GBMS has been shown, in this work and previously, to be an effective tool for evaluating and addressing such multimodality without requiring the full cost of gradient-free optimization. For future work, it is recommended that the risk of multimodality be examined in the context of multidisciplinary design optimization based on high-fidelity analysis methods, where substantial geometric freedom is almost always present.

Table 11 Degree of freedom study code reference

Test	DVs	active degrees of freedom
001_00	141	section
011_00	148	taper, section
101_00	147	twist, section
110_00	14	twist, taper
111_00	154	twist, taper, section
111_01	162	twist, taper, section, dihedral
111_11	170	twist, taper, section, sweep, dihedral

Appendix: Case Codes and Descriptions

This appendix provides a quick reference for the meaning of each test code for the degree of freedom (Table 11) and constraint (Table 12) studies. As a reminder, for the degree of freedom study, each test code is a five-digit sequence of the form *xxx_xx* where each digit corresponds to a degree of freedom and takes a zero or one to denote whether it is active, with 0 denoting the degree of freedom is inactive and 1 denoting it is active. From left to right, each digit corresponds to twist, taper, section, sweep, and dihedral. For the constraint study, each case again has a five digit code of the form *xxx_xx*, and again from left to right each digit corresponds to twist, taper, section, sweep, and dihedral, but now may take one of several values to denote which constraint (if any) is being enforced on that degree of freedom. The possible values are “L” - requiring a degree of freedom to have a linear distribution - “C” - requiring a constant value, as in a constant sweep angle - and “M” enforces a strict minimum value. The CRM geometry control system is divided into two FFD volumes, one inboard from the crank and one outboard; to permit distinct inboard and outboard distributions to develop, “L” constraints are only enforced within each FFD volume, not across them.

Table 12 Constraint study code reference

Test	Active constraints
000_00	No additional constraints
000_00-P	No additional constraint, no pitching moment
L00_00	Linear twist
0L0_00	Linear taper
00M_00	Minimum thickness
0LM_00	Linear taper, minimum thickness
LL0_00	Linear twist, linear taper
LLM_00	Linear twist, linear taper, minimum thickness
000_C0	Constant sweep
000_0C	Constant dihedral
0L0_0C	Linear taper, constant dihedral
0L0_C0	Linear taper, constant sweep
0LM_C0	Linear taper, minimum thickness, constant sweep
LL0_0C	Linear twist, linear taper, constant dihedral
LL0_C0	Linear twist, linear taper, constant sweep
0L0_CC	Linear taper, constant sweep, constant dihedral
LL0_CC	Linear twist, linear taper, constant sweep, constant dihedral
LLM_CC	Linear twist, linear taper, minimum thickness, constant sweep, constant dihedral

Acknowledgements

The authors wish to gratefully acknowledge the financial support of Bombardier Aerospace and of the Government of Ontario through the Ontario Graduate Scholarship. All cases were performed using computational resources generously provided by Compute Canada.

References

- [1] Jameson, A., “Aerodynamic design via control theory,” *Journal of Scientific Computing*, Vol. 3, No. 3, 1988, pp. 233–260, doi:<https://doi.org/10.1007/BF01061285>.
- [2] Reist, T. A., Koo, D., Zingg, D. W., Bochud, P., Castonguay, P., and Leblond, D., “Cross Validation of Aerodynamic Shape Optimization Methodologies for Aircraft Wing-Body Optimization,” *AIAA Journal*, 2020, pp. 1–15, doi:<https://doi.org/10.2514/1.J059091>.
- [3] Reist, T. A., Zingg, D. W., Rakowitz, M., Potter, G., and Banerjee, S., “Multifidelity Optimization of Hybrid Wing-Body Aircraft with Stability and Control Requirements,” *Journal of Aircraft*, Vol. 56, No. 2, 2019, pp. 442–456, doi:<https://doi.org/10.2514/1.C034703>.
- [4] Reist, T. A. and Zingg, D. W., “High-fidelity aerodynamic shape optimization of a lifting-fuselage concept for regional aircraft,” *Journal of Aircraft*, Vol. 54, No. 3, 2017, pp. 1085–1097, doi:<https://doi.org/10.2514/1.C033798>.
- [5] Chen, S., Lyu, Z., Kenway, G., and Martins, J., “Aerodynamic shape optimization of common research model wing-body-tail configuration,” *Journal of Aircraft*, Vol. 53, No. 1, 2016, pp. 276–293, doi:<https://doi.org/10.2514/1.C033328>.
- [6] Gagnon, H. and Zingg, D. W., “Aerodynamic optimization trade study of a box-wing aircraft configuration,” *Journal of Aircraft*, Vol. 53, No. 4, 2016, pp. 971–981, doi:<https://doi.org/10.2514/1.C033592>.
- [7] Méheut, M., Destarac, D., Ben Khelil, S., Carrier, G., Dumont, A., and Peter, J., “Gradient-based single and multi-points aerodynamic optimizations with the elsA software,” *53rd AIAA Aerospace Sciences Meeting*, No. 2015-0263, 2015, doi:<https://doi.org/10.2514/6.2015-0263>.
- [8] Lyu, Z. and Martins, J., “Aerodynamic design optimization studies of a blended-wing-body aircraft,” *Journal of Aircraft*, Vol. 51, No. 5, 2014, pp. 1604–1617, doi:<https://doi.org/10.2514/1.C032491>.
- [9] Ronzheimer, A., Hepperle, M., Brezillon, J., Brodersen, O., and Lieser, J., “Aerodynamic optimal engine integration at the fuselage tail of a generic business jet configuration,” *New Results in Numerical and Experimental Fluid Mechanics VIII*, Springer, 2013, pp. 25–32, doi:https://doi.org/10.1007/978-3-642-35680-3_4.
- [10] Mura, G. L., Hinchliffe, B. L., Qin, N., and Brezillon, J., “Nonconsistent mesh movement and sensitivity calculation on adjoint aerodynamic optimization,” *AIAA Journal*, Vol. 56, No. 4, 2017, pp. 1541–1553, doi:<https://doi.org/10.2514/1.J055904>.
- [11] Epstein, B., Jameson, A., Peigin, S., Roman, D., Harrison, N., and Vassberg, J., “Comparative study of three-dimensional wing drag minimization by different optimization techniques,” *Journal of Aircraft*, Vol. 46, No. 2, 2009, pp. 526–541, doi:<https://doi.org/10.2514/1.38216>.
- [12] Meheut, M., Arntz, A., and Carrier, G., “Aerodynamic shape optimizations of a blended wing body configuration for several wing planforms,” *30th AIAA Applied Aerodynamics Conference*, No. 2012-3122, 2012, doi:<https://doi.org/10.2514/6.2012-3122>.

- [13] Chau, T. and Zingg, D. W., “Aerodynamic shape optimization of a box-wing regional aircraft based on the reynolds-averaged Navier-Stokes equations,” *35th AIAA Applied Aerodynamics Conference*, No. 2017-3258, 2017, doi:<https://doi.org/10.2514/6.2017-3258>.
- [14] Long, T., Li, X., Shi, R., Liu, J., Guo, X., and Liu, L., “Gradient-Free Trust-Region-Based Adaptive Response Surface Method for Expensive Aircraft Optimization,” *AIAA Journal*, 2018, pp. 862–873, doi:<https://doi.org/10.2514/1.j054779>.
- [15] 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, doi:<https://doi.org/10.3166/remn.17.103-126>.
- [16] Chernukhin, O. and Zingg, D. W., “Multimodality and Global Optimization in Aerodynamic Design,” *AIAA Journal*, Vol. 51, No. 6, 2013, pp. 25–34, doi:<https://doi.org/10.2514/1.j051835>.
- [17] 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, doi:<https://doi.org/10.2514/1.j053318>.
- [18] 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, doi:<https://doi.org/10.2514/1.j055978>.
- [19] 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, doi:<https://doi.org/10.2514/6.2017-3752>.
- [20] Poole, D., Allen, C., and Rendall, T., “Global Optimization of Wing Aerodynamic Optimization Case Exhibiting Multimodality,” *Journal of Aircraft*, 2018, pp. 1–16, doi:<https://doi.org/10.2514/1.c034718>.
- [21] Poole, D. J., Allen, C. B., and Rendall, T., “Identifying Multiple Optima in Aerodynamic Design Spaces,” *2018 Multidisciplinary Analysis and Optimization Conference*, 2018, p. 3422, doi:<https://doi.org/10.2514/6.2018-3422>.
- [22] 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, doi:<https://doi.org/10.1016/j.ast.2018.01.016>.
- [23] Bons, N. P., He, X., Mader, C. A., and Martins, J. R., “Multimodality in aerodynamic wing design optimization,” *AIAA Journal*, Vol. 57, No. 3, 2019, pp. 1004–1018, doi:<https://doi.org/10.2514/1.j057294>.
- [24] 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, doi:<https://doi.org/10.2514/1.34810>.
- [25] 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, doi:<https://doi.org/10.2514/1.j052487>.
- [26] Hicken, J. E., *Efficient algorithms for future aircraft design: Contributions to aerodynamic shape optimization*, Ph.D. thesis, University of Toronto, 2009.

- [27] 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.
- [28] 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, doi:<https://doi.org/10.2514/1.44033>.
- [29] De Sturler, E., “Truncation strategies for optimal Krylov subspace methods,” *SIAM Journal on Numerical Analysis*, Vol. 36, No. 3, 1999, pp. 864–889, doi:<https://doi.org/10.1137/s0036142997315950>.
- [30] 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, doi:<https://doi.org/10.1137/s1052623499350013>.
- [31] Truong, A. H., Oldfield, C. A., and Zingg, D. W., “Mesh movement for a discrete-adjoint Newton-Krylov algorithm for aerodynamic optimization,” *AIAA Journal*, Vol. 46, No. 7, 2008, pp. 1695–1704, doi:<https://doi.org/10.2514/1.33836>.
- [32] 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, doi:<https://doi.org/10.2514/1.j053575>.
- [33] 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, doi:<https://doi.org/10.1145/641876.641879>.
- [34] 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, doi:[https://doi.org/10.1016/0041-5553\(79\)90085-5](https://doi.org/10.1016/0041-5553(79)90085-5).
- [35] Streuber, G. M. and Zingg, D. W., “A Parametric Study of Multimodality in Aerodynamic Shape Optimization of Wings,” *2018 Multidisciplinary Analysis and Optimization Conference*, 2018, p. 3637, doi:<https://doi.org/10.2514/6.2018-3637>.
- [36] 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, doi:<https://doi.org/10.2514/6.2015-0262>.