Investigation of multimodality in aerodynamic shape optimization based on the Reynolds-Averaged Navier-Stokes equations

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

This paper presents an improvement to a multi-start method for gradient based exploration of moderately multimodal design spaces in aerodynamic shape optimization problems. Free form deformation geometry control is used in conjunction with a gradient-based multistart algorithm to generalize cascading linear constraints for a sampler algorithm into a geometry-independent form. This permits large, high-quality, feasible samples of the design space to be generated and optimized quickly and efficiently. This algorithm is then used to explore multimodality in viscous aerodynamic shape optimization problems, including wings in transonic and subsonic flow, as well as a transonic blended wing-body design. Both wing cases are found to be multimodal, while the blended wing-body case is believed to be unimodal based on available evidence. A brief investigation of the effect that three degrees of freedom - twist, taper, and sectional control - have on multimodality in the viscous optimization of the common research model wing is undertaken. This investigation finds little evidence of multimodality related to twist and section control, while the addition of taper design variables leads to clear indications of multimodality.

I. Introduction

The key issue facing the commercial aerospace industry today is one of diminishing returns in a period of increasing demand. Unstable fuel prices and increasing public and governmental pressure to reduce greenhouse gas emissions in the face of anthropogenic climate change are driving a search for dramatically more efficient aircraft designs.[1] One of the powerful tools in this effort is aerodynamic shape optimization. Theoretically, the efficient application of large amounts of computing power to optimization could allow the discovery of novel designs not imagined by human engineers.

That potential, at the time of writing, has not yet been fully realized. While codes exist that are very effective at refining a design, they are still not capable of true, unguided invention. Overcoming this issue necessitates a thorough exploration of the design space, particularly the feasible region. Complicating these efforts is the possible presence of multimodality, i.e. problems that include multiple local optima.

The literature suggests that many common aerodynamic shape optimization problems are at least "somewhat" multimodal;[2] therefore reliable treatment of multimodality is integral to allowing one to have confidence in their exploration of the design space. Gradient-based and gradient-free algorithms offer different advantages in the face of this. Gradient-based methods[3] are relatively quick, particularly if the gradient is calculated using an adjoint method, but tend to locate local optima, and the result produced cannot be assumed to be the global optimum. Gradient-free approaches[4] are more robust than gradient-based methods with regard to filtering local optima, but can be computationally more expensive.[5] The cost of gradient-free optimization can be lessened through various methods, including hybrid algorithms,[6,7] hierarchical genetic algorithms,[8] and surrogate methods,[9–11] but is still not fully competitive with the speed of gradient-based

*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, and Associate Fellow AIAA (dwz@oddjob.utias.utoronto.ca)

optimization for many problems. Gradient-based multistart (GBMS) optimization attempts to strike a balance between speed and robustness by utilizing fast gradient-based methods in conjunction with a sampling algorithm that produces multiple starting points from the baseline geometry.[2] However, this approach alone does not address all remaining complications.

The most significant of these is how best to generate a sample of initial geometries that is geometrically plausible. Randomly assigning values to design variables is likely to produce few geometries worth investigating, but Chernukhin and Zingg[2] were able to confine sampling solely to within the feasible region by developing a system of cascading linear constraints which determine acceptable bounds on each design variable and are applied during sampling. However, due to the nature of the B-spline geometry control used, their method lacks generality, requiring specialized constraints for each class of problems and complicating application of this algorithm to new problems. By adapting Chernukhin and Zingg's algorithm for use with Free Form Deformation (FFD) geometry control, which has little dependence on the geometry being parameterized, the authors have produced a robust, geometry-independent approach for applying gradient-based methods to the exploration of multimodal design spaces. The objectives of this paper are to examine multimodality in several viscous problems, including wings under subsonic and transonic conditions, and a transonic blended wing-body (BWB), as well as to investigate the effects of several degrees of freedom on the presence and extent of multimodality in viscous optimization of the Common Research Model (CRM) wing.

II. Methodology

A. Gradient Based Multistart Algorithm - Basics

Gradient base multistart (GBMS) optimization is a term coined by Chernukhin and Zingg[2] to describe a method of utilizing the speed of gradient-based methods in the optimization of multimodal problems. The GBMS method involves the use of an initial baseline geometry and a Sobol sampling method - described in Section II - to generate a broad sample of related geometries which are then independently optimized using gradient-based methods. This attempts to converge to as many local optima as possible, thus increasing the probability that the global optimum is captured.

While the GBMS label is relatively new, equivalent methods have been used in the literature under other names.[12–16] In many of these works attention is paid to the idea of the feasible region in the design space. "Feasible" in this context goes beyond the more common definition used in most optimization applications, that is, a design for which all constraints have been satisfied. Here the "feasible region" more generally refers to the subset of all possible designs that are reasonable, meaningful, and physically realizable. This region must be carefully defined; simply randomly generating initial geometries will almost certainly produce a sample that is overwhelmingly made up of useless, physically impossible designs. However, overconstraining the initial samples could prevent a viable global optimum from being found.

A more refined approach is to create a system of linear constraints defining acceptable geometric limits on the design, which are then enforced on each sample after sampling has finished but before optimization begins, somewhat similar to the approach taken by Ugray et. al.[13] The inherent issue with this method is that enforcing linear constraints after sampling essentially snaps all non-feasible samples to the nearest boundary of the feasible region, resulting in over-sampling on the boundary of the feasible region and relatively sparse sampling within it. The ideal scenario is one where the feasible region can be fully defined a priori and sampling contained entirely within it.

In pursuit of this goal, Chernukhin and Zingg[2] developed their GBMS method using hard-coded linear constraints during the sampling process. Utilizing B-spline geometry control, their algorithm expresses upper and lower feasible limits on each design variable as a function of the values of surrounding design variables. As each design variable is sampled, the limits on subsequent variables are updated with the newly set value. This creates a cascading system of dynamic linear constraints that ensure the sampler only considers values for each design variable that lie within the user-defined feasible region.

The limitation of this method is the lack of generality in the linear constraints; each family of problems requires linear constraints tailored to that particular type of geometry. For example, the constraints from a BWB optimization cannot be applied to a planar wing case or vice-versa. Thus setup times are prohibitive and it is difficult to expand this method to new problems. Retaining the promising core of this method while reducing setup times and manual labour requirements necessitates finding a method to generalize these constraints to as many geometries as possible, ideally developing a completely geometry-independent constraint system. The solution is to adapt the B-spline based geometry control method used by Chernukhin and Zingg to a more general geometry control system, specifically FFD.

B. Free Form Deformation

FFD[17] can be visualized as embedding an object within a rubbery block which is then deformed; as the rubber block deforms the underlying object is deformed in turn. Once the underlying parameterization has been embedded in the FFD volume the interior can be treated as something of a black box from the perspective of the designer. The number, orientation, and size of the FFD volumes will change to reflect varying designs of differing complexities, but each will be a largely generic FFD volume - in this work simple three-dimensional rectangular prisms - which can contain anything from an entire wing to a single component of a complicated design. The design variables are no longer the geometry itself, or even a B-spline abstraction of it, but control points along this generic FFD volume. To ensure reasonable, feasible initial and final designs one need only constrain the relationships between these control points. These relationships are often simpler than those on the underlying geometry, but more importantly they can be programmed in such a way as to be completely independent of the embedded design. This allows a single set of highly robust linear constraints to be coded and then applied to almost any design, from a box wing to a BWB, with only minor adjustments to input parameters necessary to tailor the design space as desired.

The FFD implementation used here was originally developed by Gagnon and Zingg[18] and parameterizes the underlying geometry using B-splines. The embedded points are the control points of this B-spline parameterization, rather than the mesh nodes themselves. This ensures that at all times during deformation the original analytical definition of the shape itself remains intact.

This system was further refined via the addition of axial control.[19] This adds a NURBS curve, referred to as the "axial" curve, which is used to drive large scale deformations of the overall FFD structure. The FFD volume is divided into cross sections - planar slices of the FFD volume - which are constrained to be locally perpendicular to the axial curve. Each cross section generally includes two rows of control points, on the upper and lower surface of the FFD volume, which provide sectional control. Manipulating the axial curve through its control points deforms the entire volume, while leaving the local coordinate systems of each cross section unaffected.

With standard coordinate axes being defined as x (chord), y (span), and z (sectional or dihedral), the resulting combined system of cross sections and axial curves is defined by the degrees of freedom described below:

- Cross Sectional Degrees of Freedom
  - Twist: change in angular twist about the span axis by an entire cross section, from its initial position, measured in radians (one per cross section)
  - Taper: scaling factor of the initial chord length of each cross section (one per cross section)
  - Section: scaling factor of the initial z coordinate of each control point in a given cross section (one per cross sectional control point)
- Axial Degrees of Freedom
  - Sweep: change in global x coordinate of each axial control point from its initial position (one per axial control point)
  - Span: change in global y coordinate of each axial control point from its initial position (one per axial control point)
  - Dihedral: change in global z coordinate of each axial control point from its initial position (one per axial control point)

C. Shape Optimization Framework

Optimization is accomplished using the in-house Jetstream code,[20] which utilizes a three-dimensional, finite-difference, structured, multiblock, parallel, implicit flow solver capable of solving either the Euler[21] or Reynolds-Averaged Navier Stokes (RANS)[22] equations. The equations are discretized using second-order

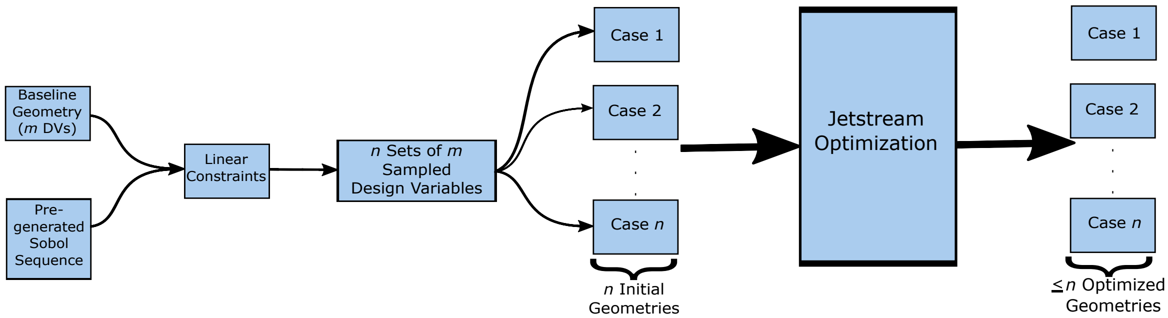

Figure 1: GBMS algorithm

summation-by-parts operators with scalar or matrix numerical dissipation, while boundary conditions and block interfaces are addressed via simultaneous approximation terms.

The foundation of the flow solver is a parallel Newton-Krylov-Schur algorithm which uses an approximate-Newton start-up phase to generate the initial iterate for the subsequent inexact-Newton phase. For both phases the resulting large system of linear equations is solved using GMRES, a preconditioned Krylov iterative solver with approximate-Schur preconditioning. For RANS analyses, the equations are closed with the Spalart-Allmaras one-equation turbulence model. Geometry parameterization and mesh movement, achieved through a linear-elasticity mesh movement algorithm, are tightly integrated.[20]

Gradients are calculated via the discrete adjoint method, a primary advantage of which is that the cost is almost independent of the number of design variables. Readers are encouraged to refer to Hicken and Zingg[20] for further details of the implementation in Jetstream. To ensure optimality, the Karush-Kuhn-Tucker (KKT) conditions must be enforced.[23] The discrete flow equations resulting from the methodology described above produce a flow adjoint system which is solved using a flexible variant[24] of the GCROT Krylov method,[25] while the solution to the mesh adjoint system is obtained using the preconditioned conjugate gradient method. The Sparse Nonlinear OPTimizer (SNOPT) accepts the calculated gradients and performs the optimization using an SQP algorithm with the quasi-Newton BFGS method used to approximate the Hessian.[26] An important feature of SNOPT in the context of this work is its ability to apply both linear and nonlinear constraints to the optimization problem, the former being critical to the FFD GBMS method developed here.

III. Gradient-Based Multistart Algorithm - Details

A flowchart of the GBMS algorithm is shown in Figure 1. The algorithm as shown is geometry control independent, that is, adapting GBMS for FFD compatibility is chiefly a task of setup modifications; beyond this point an intelligently designed GBMS code can handle almost any geometry control scheme with minor modifications.

A baseline geometry is generated beforehand; this is used as the parent of each initial geometry to follow. The sample is initially generated as a Sobol sequence - see Section A - of m real values between 0 and 1 in n dimensions, where m is the number of design variables in the problem and n is the desired number of initial geometries. Each value in this sequence is then used to scale between dynamic upper and lower limits on the corresponding design variable as follows:

$$v_{s,i} = S_{\text{low},i} + x_{d,i}(S_{\text{upp},i} - S_{\text{low},i}), \tag{1}$$

where $v_{s,i}$ is the new sampled value for the i^{th} design variable, $x_{d,i}$ is the Sobol value corresponding to the i^{th} design variable of the d^{th} sample, $S_{\text{low},i}$ and $S_{\text{upp},i}$ are the lower and upper feasible limits on this design variable, calculated using the linear constraints in Section 1 and 2. Evaluating equation (1) for each design variable results in a list of m sampled design variables, constituting a single new initial geometry. This process must then be repeated n times to produce the desired number of initial geometries. The desirable space-filling properties of Sobol sampling, combined with the dynamic nature of the linear constraints, makes it possible to produce a highly varied set of initial geometries.

Each of the obtained initial geometries is stored explicitly in a sample file as a list of modified design variable values. These are then separated into n discrete optimization problems which are run in parallel. Each initial geometry in the sample is optimized in isolation from the others and from this point the execution proceeds in the same manner as a standard gradient-based optimization problem. It is expected that some fraction of the optimizations will fail to converge for various reasons, resulting in somewhat less than n locally optimal final geometries. For this reason, separating the problems simplifies coding and improves robustness, as each design may converge, or fail to do so, without affecting the others.

In each case the optimizer is first supplied with the baseline geometry and then the sampled design variables, which it uses to distort the baseline geometry into the initial geometry. This produces a new starting geometry that is nevertheless related to the baseline design. If the sample is sufficiently large and varied, the final geometries should provide a good representation of all local optima within the design space, and the most optimal among them is likely to be the global optimum as well.

A. Sobol Sampling

The simplest approach to sampling would be to randomly assign values to each design variable; however, this may not result in an even exploration of the design space, so a more advanced sampling approach is preferred. In line with Chernukhin and Zingg,[2] this work utilizes Sobol sequence sampling to execute the algorithm outlined in Section III and evaluate the linear constraints described in Sections 1 and 2.

Sobol originally introduced these sequences as a method of approximating integrals of d-dimensional functions on a unit hypercube; however, their applications have since been extended as a robust sampling method. One of the more immediate advantages of Sobol sampling is that the number of samples, n, need not be known a priori. That is, if an initial sample of size n_1 is generated and found to be inadequate, a larger set of size n_2 can be generated without necessitating re-optimizing the first n_1 initial geometries. The specific approach utilized in this work is based on Algorithm 659,[27] which uses more primitive polynomials than other algorithms, and a Gray code implementation proposed by Antonov and Saleev[28] to generate the Sobol sequences. The Sobol sequence is generated as a preliminary step of the sampling process. Once obtained it permits the creation of the initial geometries themselves, as defined by the FFD linear constraints. These linear constraints can be logically divided into axial linear constraints - acting on the control points of the axial curves - and cross sectional linear constraints - acting on the cross sections of the FFD volumes. The following discussion is broken up in an identical fashion.

B. Linear Constraints

These linear constraints are primarily used during the sampling process, and while in some cases they, or some subset or variation of these constraints, are enforced during optimization, this is not a general rule or requirement. In the cases to follow, it will be clearly stated which linear constraints, if any, are enforced during the optimization process. The purpose of these linear constraints during sampling is to ensure that a sufficient number of the initial geometries are feasible in the sense defined in Section IIA. This improves the quality of the sample and minimizes the time the optimizer spends attempting to improve unsuitable designs. Enforcing these constraints during optimization speeds convergence by eliminating unacceptable options from the design space and reduces the appearance of infeasible features in the final designs.

1. Axial Linear Constraints

The first, and broadest, constraints acting on the axial design variables are the box constraints. These values set hard, global limits on each design variable and have a significant effect on the application of the linear constraints. The axial box constraints are defined by the user-defined vectors A_{low} and A_{upp}, each containing three values defining maximum lower and upper bounds on sweep, span, and dihedral.

Within the limits of these box constraints, the feasibility linear constraints are applied to the axial design variables in a cascading manner where the design variables at each control point are constrained in relation to the coordinates of the control point preceding it. To provide continuity in cases with multiple axial curves, the curves are connected such that the last point on a given axial curve is held to be coincident with the first point on the subsequent curve. Such a point is referred to as a "connected" axial control point. The first non-connected point in each axial curve is treated as fixed to provide an anchored reference point for the sampling process. The axial degrees of freedom were introduced previously, but it is worth reiterating here that each of the axial degrees of freedom - sweep, span, and dihedral - is respectively defined as a change in the global x, y, and z coordinate of the axial control point in question.

Due to their similar definitions, the linear feasibility constraints applied to all three design variables of the j^{th} axial control point can be expressed identically as an inequality of the form:

$$g_{j-1} + v_{j-1} + f_{j,\text{low}} \leq g_j + v_j \leq g_{j-1} + v_{j-1} + f_{j,\text{upp}}. \tag{2}$$

In the above, g represents the initial value of the coordinate being sampled at the control point in question, and v_j is the current value of the design variable being sampled. As axial design variables are defined as the change in coordinate value from the intial position, these design variables are initially zero on the baseline geometry. The variables $f_{j,\text{low}}$ and $f_{j,\text{upp}}$ are respectively the lower and upper allowable distance along a given axis between the previous axial control point $(j-1)$ and the current axial control point (j), defined as:

$$f_{j,\text{low}} = p_{j,\text{low}} C_r \quad \text{and} \quad f_{j,\text{upp}} = p_{j,\text{upp}} C_r, \tag{3}$$

where $\mathbf{p}_{\text{low}}$ and $\mathbf{p}_{\text{upp}}$ are vectors containing user defined parameters that respectively dictate feasible minimums and maximums on sweep, span and dihedral. To ensure that similar values of $\mathbf{p}_{\text{low}}$ and $\mathbf{p}_{\text{upp}}$ can be used for a variety of geometries, they are scaled by C_r, an FFD analogue of root chord calculated as the difference between the initial x coordinates of the leading edge and trailing edge FFD control points in the root cross section of the first FFD volume. Thus, in defining the values in $\mathbf{p}_{\text{low}}$ and $\mathbf{p}_{\text{upp}}$, the user is defining the maximum and minimum acceptable distance between any two adjacent axial control points, along any given axis, as a fraction of the root FFD chord. This should be relatively easy for a designer to visualize and makes the selection of parameter values more intuitive for any given case than if unscaled values were used.

For the purposes of sampling, equation (2) leads to the following S_{low} and S_{upp} values for use in equation 1:

$$\begin{aligned} S_{\text{low},j} &= g_{j-1} + v_{j-1} + f_{j,\text{low}} - g_j \\ S_{\text{upp},j} &= g_{j-1} + v_{j-1} + f_{j,\text{upp}} - g_j. \end{aligned} \tag{4}$$

2. Cross-Sectional Linear Constraints Overview

Similarly to the axial design variables, the cross-sectional design variables and all related linear constraints are bounded by a set of box constraints. Contained in the vectors B_{low} and B_{upp}, they allow the user to define a minimum and maximum acceptable value for twist, taper, and the sectional control design variables.

These box constraints set hard limits on the linear constraints, but within those bounds the freedom of motion for the cross-sectional design variables is controlled more finely by two sets of user-defined parameters: $\mathbf{r}_{\text{low}}$ and $\mathbf{r}_{\text{upp}}$, each comprising four values and respectively defining lower and upper bounds on each design variable. The parameter r_1 constrains twist, r_2 defines limits on taper, and $r_{3,4}$ handles two discrete aspects of sectional control. Unlike the axial design variables, each of the cross-sectional design variables is defined differently than the others, necessitating different constraint systems for each.

3. Twist

Each cross section has a single twist value which acts on the entire section. It is assumed that the initial twist at each cross section is zero and during optimization the overall twist at each cross section j is constrained to be within a certain range of the overall twist at the previous cross section $j-1$, defined by the values of $r_{\text{low},1}$ and $r_{\text{upp},1}$, as follows:

$$v_{j-1} + r_{\text{low},1} \leq v_j \leq v_{j-1} + r_{\text{upp},1}. \tag{5}$$

During sampling, the upper and lower sampling limits from equation (1) are obtained by separating this inequality into two equations:

$$\begin{aligned} S_{\text{low},j} &= v_{j-1} + r_{\text{low},1} \\ S_{\text{upp},j} &= v_{j-1} + r_{\text{upp},1}. \end{aligned} \tag{6}$$

4. *Taper*

Similarly to twist, each cross section has a single taper value. The design variable itself is a scaling factor on the initial FFD chord - the difference between local x coordinates of the leading and trailing edge FFD control points. Thus a taper value of 2 implies the chord has doubled during optimization, a value of 0.5 implies the chord has halved. The linear constraint does not constrain the design variable directly, but permits the chord at the j^{th} cross section to only take values within an acceptable range of the chord at the $j-1$ cross section, as defined by:

$$C_{j-1}v_{j-1} + r_{\text{low},2}C_r \leq C_j v_j \leq C_{j-1}v_{j-1} + r_{\text{upp},2}C_r. \tag{7}$$

Here C_j represents the initial FFD chord at each cross section and the FFD root chord, C_r, is merely $C_{j=1}$. Once again, breaking apart the inequality yields the equations used during sampling to obtain the upper and lower sampling limits:

$$\begin{aligned} S_{\text{low},j} &= \frac{C_{j-1}v_{j-1} + r_{\text{low},2}C_r}{C_j} \\ S_{\text{upp},j} &= \frac{C_{j-1}v_{j-1} + r_{\text{upp},2}C_r}{C_j}. \end{aligned} \tag{8}$$

5. *Sectional Control*

The sectional control represent the most challenging design variables to constrain in this system due to their unique properties. While twist and taper act on an entire cross section, each control point within the cross section has its own z coordinate controlled by its own design variable. Each of these sectional design variables are a separate scaling factor on the initial local z coordinate of each control point within the cross section. For this reason any points with initial z coordinates of 0 should be avoided, though this can be a useful tool to selectively fix certain control points along the z axis if desired. Constraining this system fully requires two related but discrete linear constraints.

The first sectional constraint acts on the control points along the lower surface of the FFD volume and bounds the z coordinate of each to be within a certain range of the z coordinate of the previous point, that is, the adjacent point closer to the leading edge, preventing excessively jagged airfoil cross sections. In each cross section the first point along the lower surface is constrained relative to the corresponding point on the previous cross section, moving inboard to outboard, the exception being the first point of the first cross section, which is held constant to provide a fixed reference point during sampling. The range of acceptable values is defined by $r_{\text{low},3}$ and $r_{\text{upp},3}$ and is normalized by the root FFD chord in a similar fashion as the taper constraint.

Other than the aforementioned exception for the first point in each cross section, the general linear constraint for the j^{th} control point along the lower surface of each cross section can be written as:

$$g_{z,ls,j-1}v_{ls,j-1} + r_{\text{low},3}C_r \leq g_{z,ls,j}v_{ls,j} \leq g_{z,ls,j-1}v_{ls,j-1} + r_{\text{upp},3}C_r, \tag{9}$$

where $g_{z,ls,j}$ denotes the initial local z coordinate of the j^{th} point along the lower surface of the cross section in question in the local coordinate system of that cross section. This yields the sectional sampling constraints for the lower surface control points as follows:

$$\begin{aligned} S_{\text{low},ls,j} &= \frac{r_{\text{low},3}C_r + g_{z,ls,j-1}v_{ls,j-1}}{g_{z,ls,j}} \\ S_{\text{upp},ls,j} &= \frac{r_{\text{upp},3}C_r + g_{z,ls,j-1}v_{ls,j-1}}{g_{z,ls,j}}. \end{aligned} \tag{10}$$

The second linear sectional constraint maintains a minimum and maximum acceptable cross section thickness, defined by $r_{\text{low},4}$ and $r_{\text{upp},4}$ and normalized by the root FFD chord C_r, by constraining each point along the upper surface to the corresponding lower surface control point. The resulting linear constraint between the j^{th} point along the upper surface of a cross section and the corresponding j^{th} control point on the lower surface of the same cross section is:

$$r_{\text{low},4}C_r \leq g_{z,us,j}v_{us,j} - g_{z,ls,j}v_{ls,j} \leq r_{\text{upp},4}C_r, \tag{11}$$

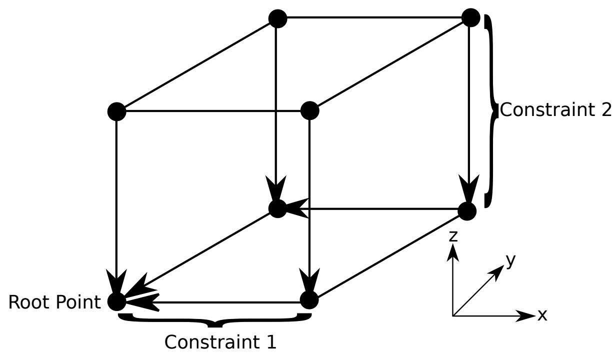

Figure 2: Sectional constraint illustration

where $g_{z,ls,j}$ is defined identically as in equation (9), and $g_{z,us,j}$ is similarly the initial z coordinate of the j^{th} point along the upper surface of the same cross section. During sampling, upper and lower sampling constraints are calculated for each point along the upper surface according to

$$\begin{aligned} S_{\text{low},us,j} &= \frac{r_{\text{low},4} C_r + g_{z,ls,j} v_{ls,j}}{g_{z,us,j}} \\ S_{\text{upp},us,j} &= \frac{r_{\text{upp},4} C_r + g_{z,ls,j} v_{ls,j}}{g_{z,us,j}}. \end{aligned} \tag{12}$$

To illustrate the interaction of these linear constraints, Figure 2 shows the sectional constraint system for a simple two cross section FFD volume; for clarity the axial curve is omitted. Each arrow begins at the point that is constrained and ends at the point it is constrained relative to. Following the arrows, one can trace a path of cascading constraints connecting any point to the root point, which acts as a fixed reference point during sampling.

IV. Results

Four cases are examined: exploratory optimization of a planar wing in subsonic and transonic flow, a more constrained optimization of a BWB in transonic flow, and a study of the effects of various degrees of freedom on multimodality using the common research model (CRM)[29] wing under transonic conditions. In the following discussions, several mentions are made of the concept of a targeted number of "viable" samples, which refer to initial geometries that do not fail on the initial flow solution or mesh movement. These are important because it is more efficient for a sample to fail early in the optimization process than later in the process once potentially significant computational resources have been expended attempting to improve the failed design. Ideally, failures should be minimized to improve efficiency, without being eliminated entirely, to ensure that the problem is not overly constrained. Generally speaking a failure rate of 10% was targeted, and early failures were favoured over later failures. However, this was not necessarily always possible to achieve. To test for convergence of the optimization problems, three primary metrics were used. The first is feasibility, which in this context can be taken as a measure of the nonlinear constraint violations, the second is optimality, which is the gradient of the merit function - the objective function augmented with terms for the linear and nonlinear constraints - and the third is the relative merit function, which is the ratio of the merit function on the first iteration, to that on the current iteration.

A final note on methodology is worthwhile: in all cases it was found that the apparent degree of multimodality was heavily dependent on the degree of convergence achieved, with tighter convergence being associated with lower degrees of multimodality. To improve performance and reduce computational cost, a screening mechanism was employed. All initial geometries were optimized, if possible, until feasibility was satisfied and a noticeable drop in optimality achieved. At this point, any so-called "intermediate" optima -

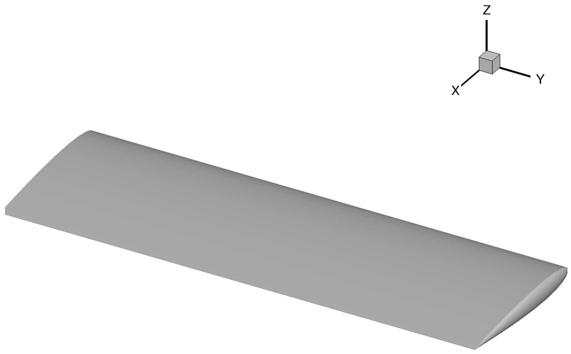

(a) Baseline geometry

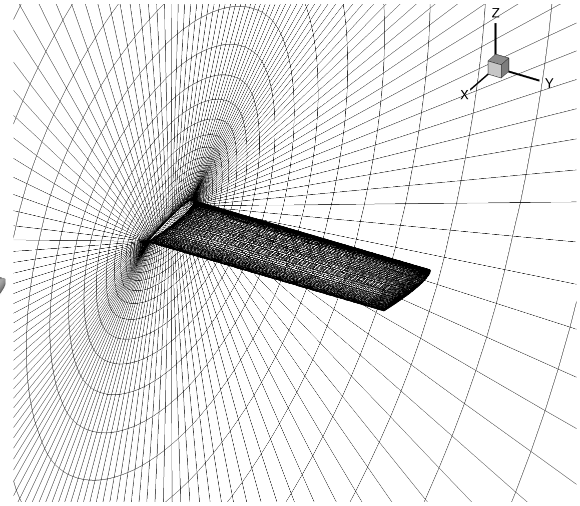

(b) Baseline optimization mesh

Figure 3: Baseline wing case

geometries that one or more initial geometries had converged to - were noted. It was assumed from this point that if any two initial geometries had reached the same intermediate optimum they would follow a similar path during subsequent optimization, thus only a single example of each intermediate optimum was carried forward and optimized further until optimization stalled. The resulting final optima and their performance values are those that are ultimately reported here. It was common in most cases for two or more intermediate optima to converge to the same, or highly similar, final optima and it is assumed that any two examples of a given intermediate optimum would converge to the same final optimum given sufficient optimization time.

A. Exploratory Wing Optimization Under Subsonic and Transonic Conditions

The wing optimization cases are based on a viscous transonic exploratory planform optimization case performed by Koo and Zingg,[30] expanded to include both subsonic and transonic flows. The baseline geometry for both cases, shown in Figure 3(a), has a rectangular planform, the NACA 0012 airfoil, and is normalized by initial root chord, with an initial semi-span of 2.92, projected area of 2.92, and initial volume of approximately 0.250. The design variables at each cross section are twist, taper and sectional control, with twist at the root cross section being fixed in favour of an angle-of-attack (AOA) design variable. Large-scale deformations are driven by the axial curve design variables sweep, span and dihedral, all of which are fixed at the root axial control point. The feasibility constraints presented previously are enforced during both sampling and optimization in these cases, with the values assigned to these parameters listed in Table 1.

Applying the sampling algorithm to the baseline geometry in Figure 3a, utilizing the parameters in Table 1, 400 initial geometries were generated, a sample of which - including the baseline - is shown in Figure 4. It is apparent from Figure 4 that the algorithm is capable of producing numerous diverse geometries from a highly basic baseline geometry, while nevertheless still constraining itself to generally reasonable shapes.

1. *Optimization Problem*

The optimization problem is lift-constrained drag minimization at a freestream Mach number of 0.5 for the subsonic case and a Mach number of 0.8 for the transonic case. Several nonlinear constraints are also enforced: the maximum root bending moment is set as 80% of the bending moment produced by the baseline geometry optimized without a bending moment constraint, the coefficient of lift, C_L, is constrained to 0.2625, projected area to 3.06, and volume to no less than 0.245, normalized by the root chord of the baseline geometry. The projected area and volume values differ slightly from the projected area and volume of the baseline geometry, the constraint values arising from an earlier inviscid GBMS optimization of a planar NACA 0012 baseline geometry, for which the initial volume and projected areas matched the constraints. This discrepancy between the inviscid and viscous baseline geometries is produced by a change in the type of

Table 1: Linear constraint parameter values for wing optimization cases

Parameter	Value
A_{low}	$(-1.00, -0.60, -0.45)$
A_{upp}	$(1.00, 0.60, 0.45)$
B_{low}	$(-0.0545, 0.45, 0.50)$
B_{upp}	$(0.0545, 1.55, 1.50)$
p_{low}	$(-0.50, 0.50, -0.30)$
p_{upp}	$(0.50, 6.00, 0.30)$
r_{low}	$(-0.0349, \frac{-11}{30}, 0.025, -0.100)$
r_{upp}	$(0.0349, \frac{11}{30}, 0.500, 0.100)$
α_{low}	-3.0^{o}
α_{upp}	6.0^{o}

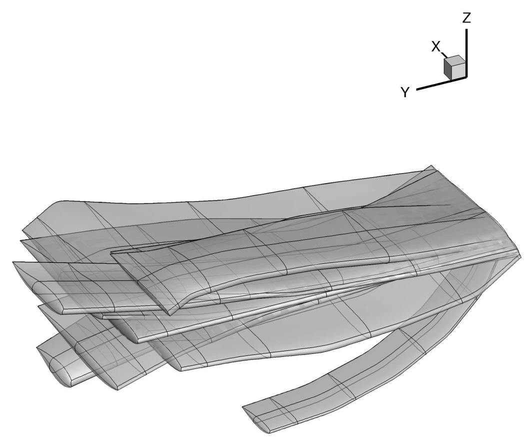

Figure 4: Selection of initial geometries from NACA 0012 sample

wingtip "cap" that is used at the end of the geometry; the inviscid geometry utilized a pinched tip, while the viscous cases have a more rounded wingtip geometry. This decision was made simply to better accommodate the different meshes used for the two sets of cases and should not significantly affect the results.

The problem can be summarized as

$$\begin{aligned} \text{minimize} \quad & C_D \\ \text{w.r.t.} \quad & v \\ \text{subject to} \quad & C_{\text{lin}} \geq 0 \\ & C_M b \leq 0.3271 \\ & C_L = 0.2625 \\ & S = 3.06 \\ & V \geq 0.245 \end{aligned} \tag{13}$$

where C_D is the coefficient of drag, C_L is the coefficient of lift, C_M is the root bending moment coefficient, b is the semispan, S is the projected area, V is the volume, v are the geometric design variables, and C_{lin} are the linear constraints.

2. *Optimization Grid*

The optimization grid comprises an O-O topology consisting of 64 blocks and 1,320,000 nodes, shown in Figure 3b. Fine-mesh results were obtained using Richardson extrapolation,[31] which uses results from increasingly fine meshes to extrapolate to continuum values on a theoretical infinitely fine mesh. Including the optimization mesh, three mesh levels were used for this process - a selection of important values for each mesh are listed in Table 2.

Table 2: Wing mesh parameters

Grid	Nodes	Off-wall Spacing (c)	Average y^+
L0	1,320,000	1.3×10^{-6}	0.59
L1	2,521,664	1.0×10^{-6}	0.44
L2	4,964,544	7.9×10^{-7}	0.32

3. *Geometry Control*

Geometry control for the wing optimization cases, both subsonic and transonic, is accomplished through six cross sections which provide sectional control, with large deformations driven by a single axial curve controlled by four equispaced axial control points. Each cross section is allowed twist, taper, and sectional degrees of freedom, with the exception of the root cross section, for which the twist is held fixed in favour of an AOA design variable. Section control is provided by 10 section design variables per cross section, or 60 total. The root axial control point is held fixed in all three dimensions, changes in sweep, span, and dihedral are permitted for the three remaining axial control points.

4. *Subsonic Wing Optimization*

Optimization of the subsonic case was performed on the coarse optimization mesh of 1,320,000 nodes at a Reynolds number of 12,500,000 and a Mach number of 0.50. From the 400 generated initial geometries, 61 "viable" samples - initial geometries that do not fail on the first flow solve or mesh movement - were optimized; this represents 60 samples plus the baseline geometry. This threshold was met after attempting 65 initial geometries, four of which failed immediately; however, seven of the remaining 61 samples subsequently failed to converge, leaving 54 optimized final geometries and an ultimate success rate of approximately 84%. Eleven locally optimal geometries were obtained, all of which are shown in Figure 5. Figure 6, which displays the convergence histories for each optimum, indicates that acceptable convergence was achieved and the feasibility tolerance of 10^{-6} was satisfied for each optimum.

Performance results for each local optimum and the unoptimized baseline are tabulated in Table 3. The ratio L/D is the preferred performance metric in this case, as a result of small variations in the C_L values for each optimum which appeared when the converged solutions were re-evaluated on the subsequent finer meshes. The "number of examples" column indicates the number of initial geometries that converged to that particular optimum, or to intermediate optima that were found to later converge to that particular optimum. The optimized baseline here was found to be one of the better local optima, being outperformed by just one other design. There is also a tendency for the optima to cluster into groups of similarly performing geometries, for instance optimum 1, 2, 6, and 7 offer almost identical performance despite being clearly geometrically differentiated. Nevertheless, the L/D values of the various local optima differ by up to 3.38% of the optimized baseline value. It is also worth noting that the best geometry is not the most, or even second most, common local optimum. This reinforces the importance of a sufficient exploration of the design space, as a single random sample or half-hearted exploration would be unlikely to locate the global optimum in this case.

5. *Transonic Wing Optimization*

The transonic wing optimization case is identical to the subsonic case, save that it is optimized at a Mach number of 0.80 and a Reynolds number of 20 million. Once again, 61 viable samples were located and

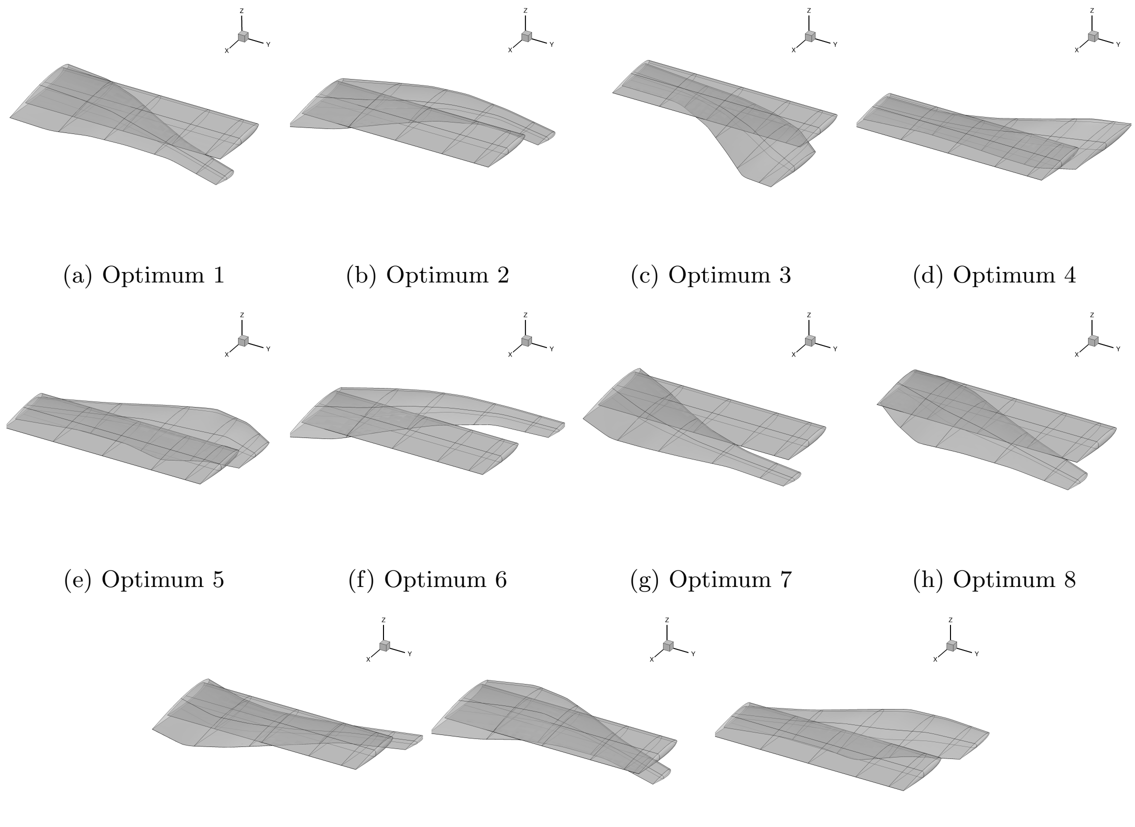

(a) Optimum 1 (b) Optimum 2 (c) Optimum 3 (d) Optimum 4

(e) Optimum 5 (f) Optimum 6 (g) Optimum 7 (h) Optimum 8

(i) Optimum 9 (j) Optimum 10 (k) Optimum 11

Figure 5: Subsonic wing, local optima compared against baseline geometry

optimized as fully as possible. Locating these samples proved more difficult than for the subsonic case, 126 initial geometries were attempted before 61 could be located that did not fail on the initial flow solve. Further, 10 of the 61 viable samples failed to satisfy the feasibility constraint within a generous time limit and were therefore concluded to have failed, resulting in 51 optimized geometries and a success rate of just 40%. To allow comparison between this case and the subsonic problem, the linear constraints were kept as is. Furthermore, the nature of the Sobol sampling method used provides confidence that a thorough exploration of the design space was achieved. Nevertheless, this strongly suggests that in future applications efficiency can be improved by tightening the linear constraints for transonic RANS problems.

The 55 successful samples led to eight distinct local optima, i.e. the optimization problem is “somewhat” multimodal according to Chernukhin and Zingg,[2] as shown in Figure 7. These results are geometrically very similar to those yielded by the subsonic case, with optima 1, 2, 4, 6, and 7 all having very similar counterparts from the subsonic optima set - Figure 5. The two sets of optima are not identical, however, and the differences evince the features one expects in a transonic optimization. The transonic optima generally incorporate more non-planar features and a greater degree of sweep than comparable subsonic optima; further, a strong preference for backswept wings is noted, with forward-swept geometries comprising 55% of the subsonic optima and just 25% of the transonic optima.

Examining the tabulated performance values in Table 4, several points become apparent. First the transonic test offers much greater potential for performance improvement from the unoptimized baseline than the subsonic test. However, this is wholly predictable and more attributable to the obvious unsuitability of the baseline geometry for transonic flight than any increased effectiveness of the optimizer under transonic conditions. The transonic optima produce a slightly smaller range of performances than the subsonic case, but the presence of multimodality is by no means negligible, representing up to a 2.56% difference in performance depending on which initial geometry is used. This serves as strong evidence that multimodality is both

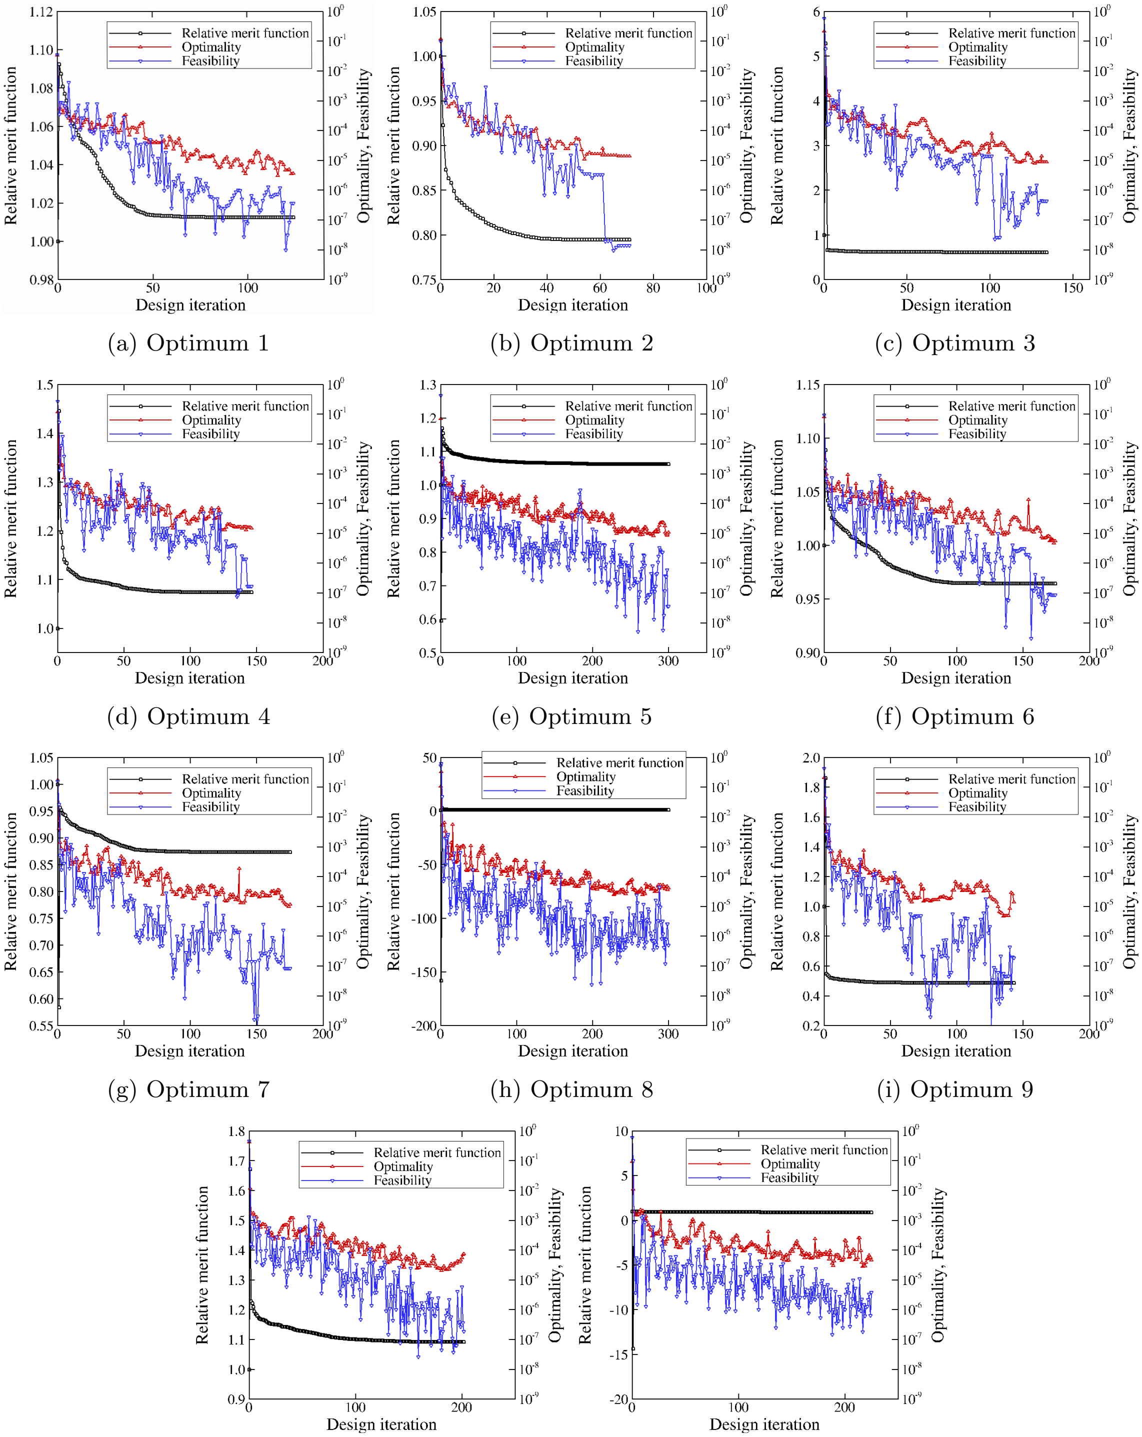

(a) Optimum 1 (b) Optimum 2 (c) Optimum 3

(d) Optimum 4 (e) Optimum 5 (f) Optimum 6

(g) Optimum 7 (h) Optimum 8 (i) Optimum 9

(j) Optimum 10 (k) Optimum 11

Figure 6: Subsonic wing local optima convergence histories

Table 3: Subsonic wing performance results, Richardson extrapolated values

Value	C_L	C_D	C_L/C_D	% change from optimized baseline	Number of examples
Baseline	0.2788	1.162×10^{-2}	23.992	-	-
Optimized Baseline (1)	0.2663	1.033×10^{-2}	25.776	0.00%	16
Optimum 2	**0.2664**	$\mathbf{1.027 \times 10^{-2}}$	**25.929**	**0.59%**	**10**
Optimum 3	0.2675	1.055×10^{-2}	25.358	-1.62%	11
Optimum 4	0.2671	1.062×10^{-2}	25.150	-2.43%	1
Optimum 5	0.2662	1.062×10^{-2}	25.078	-2.71%	1
Optimum 6	0.2674	1.038×10^{-2}	25.763	-0.05%	8
Optimum 7	0.2661	1.033×10^{-2}	25.762	-0.05%	3
Optimum 8	0.2659	1.050×10^{-2}	25.315	-1.79%	1
Optimum 9	0.2672	1.039×10^{-2}	25.733	-0.17%	1
Optimum 10	0.2662	1.036×10^{-2}	25.704	-0.28%	1
Optimum 11	0.2664	1.063×10^{-2}	25.056	-2.79%	1

Table 4: Transonic wing performance results, Richardson extrapolated values

Value	C_L	C_D	C_L/C_D	% change from optimized baseline	Number of examples
Baseline	0.2819	2.586×10^{-2}	10.901	-	-
Optimized Baseline (1)	0.2676	1.031×10^{-2}	25.963	0.00%	19
Optimum 2	0.2679	1.034×10^{-2}	25.917	-0.18%	6
Optimum 3	0.2690	1.053×10^{-2}	25.541	-1.63%	4
Optimum 4	**0.2695**	$\mathbf{1.038 \times 10^{-2}}$	**25.966**	**0.01%**	**12**
Optimum 5	0.2683	1.061×10^{-2}	25.301	-2.55%	1
Optimum 6	0.2693	1.046×10^{-2}	25.747	-0.83%	5
Optimum 7	0.2683	1.050×10^{-2}	25.548	-1.60%	3
Optimum 8	0.2698	1.047×10^{-2}	25.497	-0.71%	1

present and potentially significant in RANS optimization. Once again we find that the most common local optimum is not the best; however, in this case the difference between the best and most numerous optima is largely negligible and together they account for a significant majority of the final geometries.

B. Transonic Blended Wing Body

The BWB examined here is an FFD-controlled case previously presented by Lee et al.,[32] itself adapted from a B-spline-controlled case presented by Reist and Zingg.[33] The baseline geometry has an initial span of 1.46 and a root chord of 1.68, with dimensions scaled by the mean aerodynamic chord; the geometry is shown in Figure 8(a). The cross sectional design variables are twist, taper and section, with the root twist fixed in favour of an angle-of-attack design variable. Axial degrees of freedom are permitted as well, specifically sweep and the span of the wing and body; however total span is fixed.

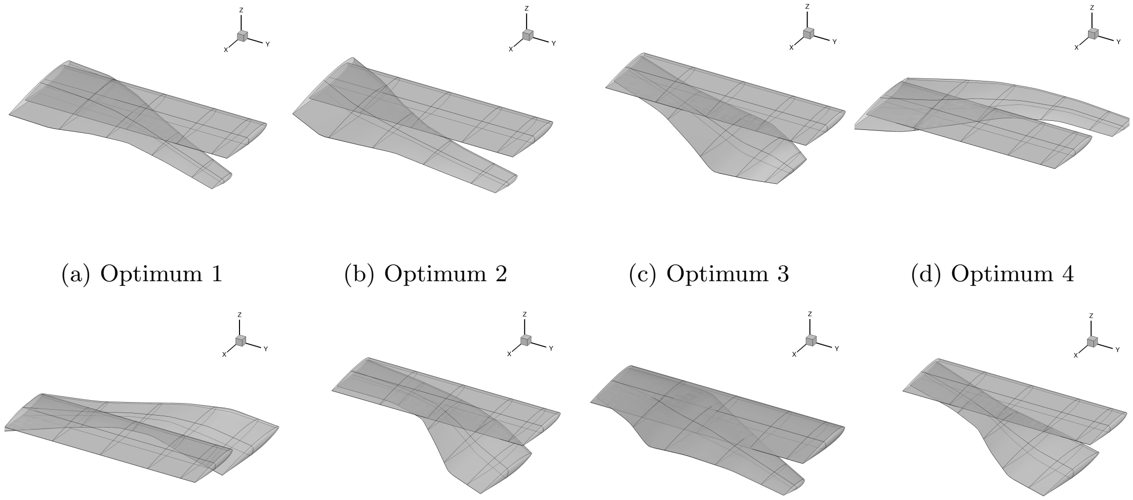

(a) Optimum 1 (b) Optimum 2 (c) Optimum 3 (d) Optimum 4

(e) Optimum 5 (f) Optimum 6 (g) Optimum 7 (h) Optimum 8

Figure 7: Transonic wing, local optima compared against baseline geometry

1. Optimization Problem

The objective is to maximize the lift-to-drag ratio at a Mach number of 0.8 and a Reynolds number of 62 million. No C_L constraint is enforced during optimization. During optimization the volume of the body - excluding the wing sections - is constrained to be no less than 0.0744 and the projected area of the body is fixed at 0.466, while the volume of the wing sections must be at least 0.00558, corresponding to the volumes and projected areas of the baseline geometry. These values represent a regional-class BWB with a total span of 118 feet and a passenger capacity equivalent to that of an Embraer E190. The problem can be written formally as

$$\begin{aligned} \text{maximize} \quad & \frac{C_L}{C_D} \\ \text{w.r.t.} \quad & v \\ \text{subject to} \quad & C_{\text{lin}} \geq 0 \\ & S_{\text{body}} = 0.466 \\ & V_{\text{body}} \geq 0.0744 \\ & V_{\text{wing}} \geq 0.00558 \end{aligned} \tag{14}$$

S_{body} is the projected area of the body, and V_{body} V_{wing} are respectively the volume of the body and wing sections.

Parameters used for the developed linear feasibility constraints can be found in Table 5; as with the wing cases these constraints are applied during both sampling and optimization. In addition, a set of linear constraints is applied to maintain a general BWB shape in all sampled geometries. Linear twist is enforced on the body and wing sections, linear taper is maintained on the wing, while taper on the body is allowed to deviate up to 25% from linear. The ratio of tip to root taper on each axial must be between 1 and 0.1. Additionally, the sweep angle is fixed at the initial value. These constraints, along with the linear feasibility constraints are applied during both sampling and optimization in this case.

2. Optimization Grid

The optimization mesh, denoted as L0, shown in Figure 8(b), and with key parameters as per Table 6, is an H-H grid consisting of 1,225,728 nodes. Similarly to previous cases, this coarse mesh was used to locate

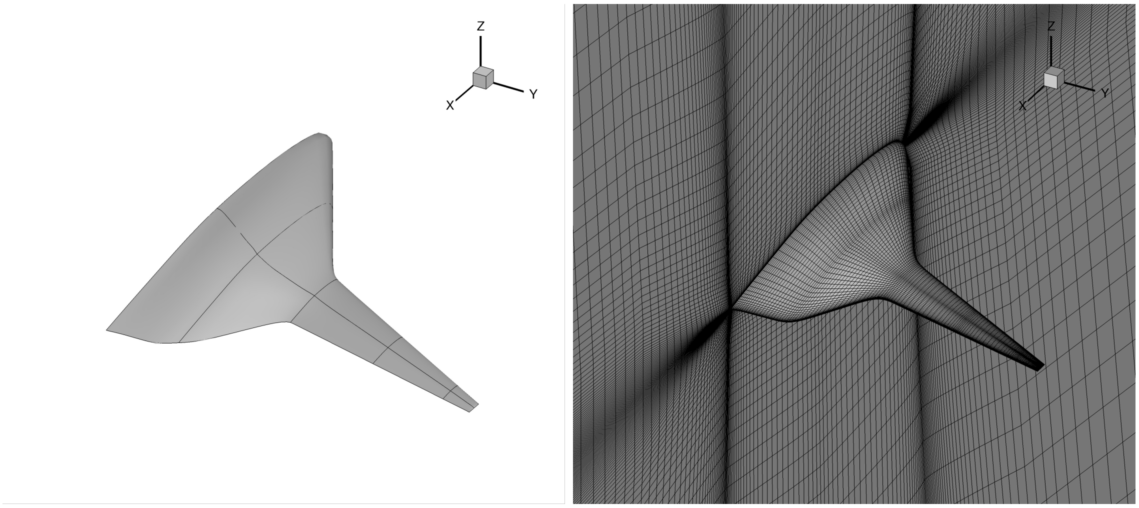

(a) Baseline geometry (b) Baseline optimization mesh

Figure 8: Baseline BWB design

Table 5: Linear constraint parameter values for BWB case

Parameter	Value
A_{low}	$(-1.00, -0.675, -10.0)$
A_{upp}	$(2.5, 0.675, 10.0)$
B_{low}	$(-0.1745, 0.5, 0.25)$
B_{upp}	$(0.1745, 1.5, 2.00)$
p_{low}	$(-0.050, 0.01, -0.25)$
p_{upp}	$(0.175, 0.75, 0.25)$
r_{low}	$(-0.0174, -0.18, 0.0025, -0.0125)$
r_{upp}	$(0.0174, 0.05, 0.25, 0.0140)$
α_{low}	-3.0^{o}
α_{upp}	3.0^{o}

any local optima, which were then re-solved on progressively finer meshes (L1 and L2), with Richardson extrapolation used to estimate grid converged results.

3. BWB Geometry Control

Unlike the other cases examined here, two connected FFD volumes are used, one controlling the aircraft body and the other the wing. The volumes are divided into 7 and 11 cross sections respectively, with section control for each supplied by 22 control points per cross section. Large-scale deformations of each design variable are driven by two axial curves, each associated with one of the FFD volumes and controlled by 20 axial control points. To ensure continuity, the axial curves are connected such that the "tip" control point of the body axial curve is coincident with the "root" control point of the wing axial. All degrees of freedom are active in each cross section - twist, taper, and section - with the exception of the root cross section of the body, for which twist is kept fixed in favour of AOA which is instead used as a design variable. Sweep and span changes are permitted at all axial control points, save the root control point of the body, which is fixed, and the tip control point of the wing which is only allowed to vary sweep. Dihedral is fixed at zero for all axial control points.

Table 6: Blended wing-body mesh parameters

Grid	Nodes	Off-wall Spacing (mean aerodynamic chord)	Average y^+
L0	1,225,728	1.3×10^{-6}	2.13
L1	2,374,848	9.5×10^{-7}	1.56
L2	5,077,800	7.2×10^{-7}	1.19

4. BWB Optimization

A sample of 500 initial geometries was generated by the sampling algorithm, a selection of which are shown in Figure 9. This showcases the variety of geometries that the sampler is able to generate even within the constraints imposed by the problem definition. It should also be reiterated that the linear feasibility constraints applied here are identical to those applied in the wing optimization cases, differing only in the selection of control parameters. The additional linear constraints mentioned above are not necessary to generate samples of this geometry, but were included to further constrain sampling to a more specific subset of the feasible region.

A deeper investigation of this case was undertaken, with 101 viable geometries being targeted, including the baseline. This proved to be a computationally difficult problem, requiring 290 attempted samples to locate 101 viable initial geometries. Furthermore, 33 of these 101 geometries failed to converge in any meaningful sense, producing 68 final geometries and a success rate of just 23%. Efforts to improve this success rate through further tuning of the parameters in the linear constraints proved largely ineffective as there was no clear monotonic relationship in this case between success rate and the tightness of the linear constraints. This may simply be due to the computational difficulty of this problem. The 68 final geometries represent an exploration at least on par with that performed in the other cases and are sufficient to make similar conclusions about the characteristics of the design space.

Forty-eight of the 68 initial geometries, approximately 70%, including the baseline geometry, converged to precisely the same local optimum. Of the 20 remaining initial geometries, no two converged to the same final design, with a significant spread of performance and geometric features being apparent in the results, though all are significantly inferior in performance to the "primary" local optimum located by the majority of initial geometries. A more detailed examination of these final geometries reveals that the 21 local optima, including the primary optimum, are not wholly unique, but can be categorized into four families based on largely similar geometric features, but often still displaying significant variations in design and performance. The final geometries for each of these four families are shown in Figure 10.

Figure 11 displays the convergence histories for one example case from each family. It is clear that one of the differentiating features between Family 1 and Families 2 through 4 is the degree of convergence achieved. While Family 1 shows clear convergence, typified by satisfaction of the feasibility tolerance, a clear drop in optimality, and asymptotic behaviour of the relative merit function, the remaining families consistently fail to achieve the first and last requirements, even after extended optimization. The implications this has for our conclusions regarding the design space will be discussed in greater detail later.

Estimated grid converged results for each family are shown in Table 7; these were obtained by evaluating the flow for a single member of each family on the finer grids. With the exception of family 1, a notable degree of variation in performance is still apparent within each family, therefore, these values are not conclusive or definitive statements on the performance of each family, but provide a rough idea of the general spread of performance that was found throughout the final geometries. Given the uniformly poor performance of the geometries in families 2-4, and consideration of further evidence given below, it is believed that this problem is in fact unimodal and the additional families represent only partially optimized designs which have stalled due to various difficulties during optimization. This is corroborated by the convergence data given in Figure 11, which shows that while the feasibility tolerance was satisfied and merit function reductions were achieved in every case, for families 2-4 the optimizer struggles to produce any clear reduction in optimality. The poor - or nonexistent - reduction in optimality for families 2-4 suggests that the optimizer became stuck before ultimately failing. These failures were primarily caused by repeated failed flow evaluations or mesh movements, eventually leading to the optimizer exhausting available search directions to move forward. This could be due to any one of a number of reasons, ranging from the more highly constrained nature of this

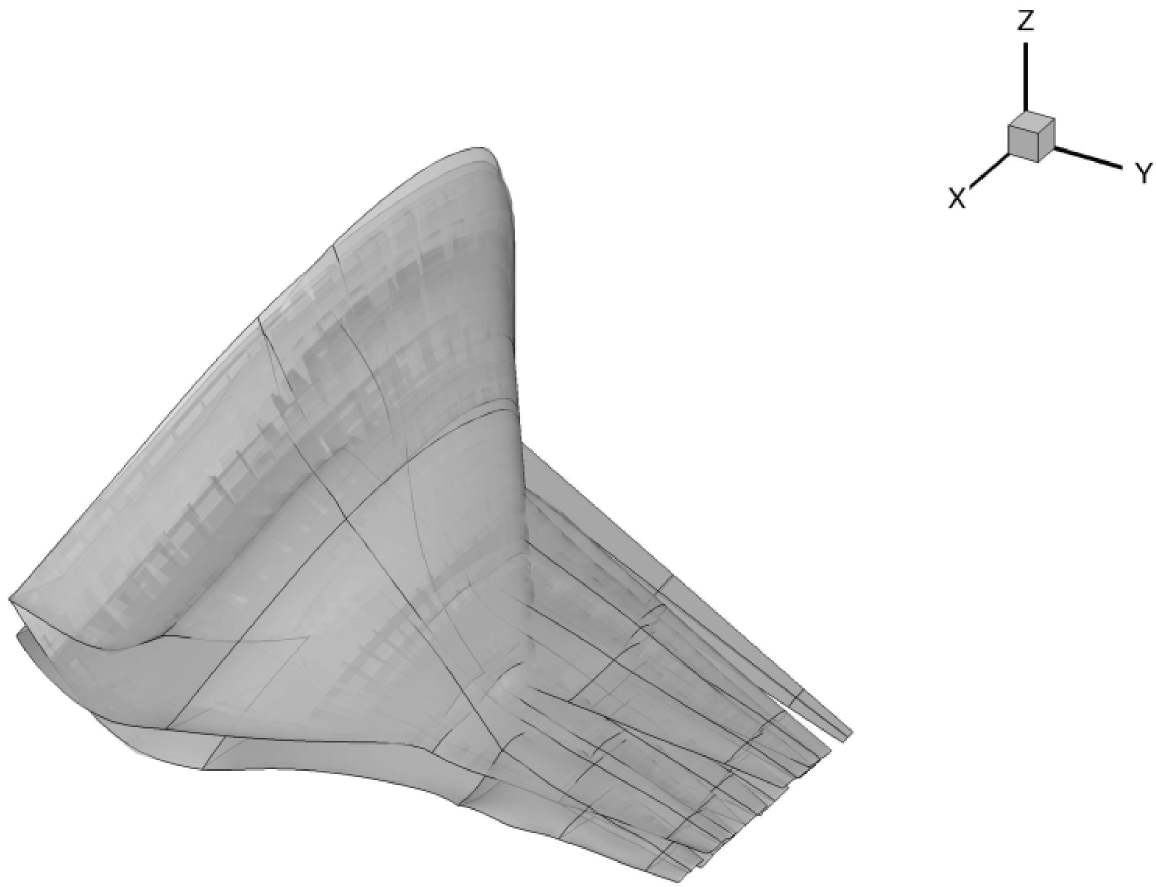

Figure 9: Selection of initial geometries from BWB sample

problem to the design space simply proving difficult to navigate. It should be noted for clarity that any gaps in the feasibility plots represent regions where the the feasibility was exactly zero, i.e. the nonlinear constraints were satisfied exactly.

The data on the number of final geometries that can be categorized into each family is also valuable, and reinforces the unique nature of family 1. No other family had more than approximately 20% as many members as family 1, and the second most common group was in fact failed cases (33 examples) which still accounted for only 70% as many cases as the first family. Further, it should be reinforced that while families 2-4 represent a variety of geometries having similar major characteristics, family 1 is entirely composed of cases that have converged to the same, nearly identical final geometry. This strongly suggests that there is only one true optimum in this design space, or at minimum that the global optimum is highly dominant over any other local optima. This also reinforces the necessity of thorough convergence when performing optimization in multimodal design spaces; insufficient convergence of the geometries in this case would have presented results apparently indicating an extremely multimodal design space without any clear global optimum. Good judgment on the part of the designer is also critical, as the decision on what constitutes a new "local optimum" versus a less optimized relative of a given optimum, or how to define sufficient convergence, will always be subjective.

This conclusion suggests an explanation for the unusual "familial" clustering of the 20 outlying final

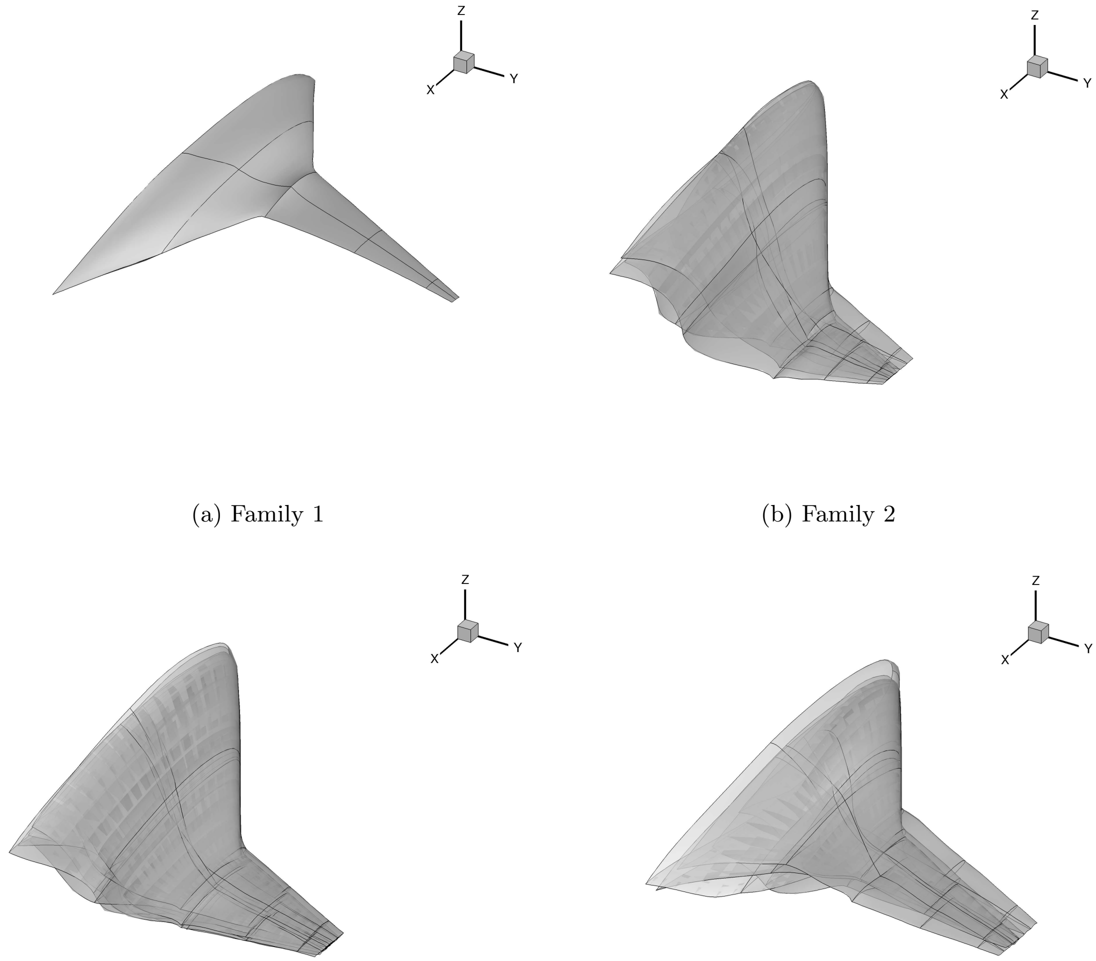

(a) Family 1

(b) Family 2

(c) Family 3

(d) Family 4

Figure 10: Geometry families produced by BWB case

geometries. It was earlier noted that during optimization the cases tended to converge to a larger set of “intermediate” local optima, which on further optimization would ultimately converge to a smaller set of final local optima. The families of geometries located here may represent a rough depiction of such intermediate local optima, that if the optimizer had been able to proceed would have eventually been driven on to the primary global optimum. To reinforce the point made earlier, in every single case in which feasibility was satisfied and significant optimality reductions achieved, the geometry in question invariably converged to the same final optimum. The only geometries which did not converge to this global optimum were those that failed to adequately optimize the initial geometry. It is the belief of the authors that the balance of evidence suggests this case is unimodal. However, it is impossible to rule out the possibility of multimodality for viscous optimization of BWB geometries with a greater degree of geometric flexibility.

C. Common Research Model Multimodality Study Preliminary Results

This case examines an FFD parameterization of the Common Research Model (CRM) wing used by Lee et al.[32] The baseline geometry, shown in Figure 12(a) is derived from the CRM wing-body configuration used in the Fifth Drag Prediction Workshop[29] with the fuselage removed and the geometry scaled by the mean aerodynamic chord. Optimization occurs at a Mach number of 0.85 and a Reynolds number of 5

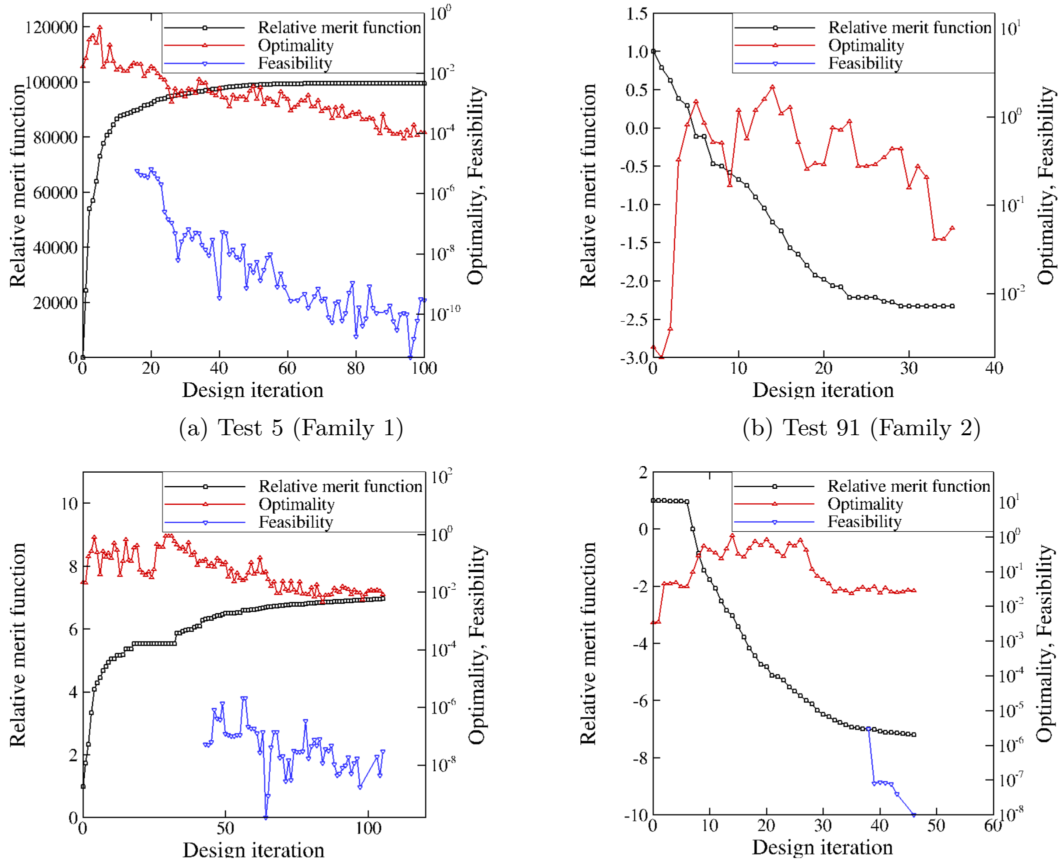

(a) Test 5 (Family 1)

(b) Test 91 (Family 2)

(c) Test 51 (Family 3)

(d) Test 62 (Family 4)

Figure 11: Blended wing-body, selection of convergence histories. Any gaps in the feasibility plots represent regions where the feasibility was exactly zero, i.e. the nonlinear constraints were satisfied exactly; for Figure 11(b) this extends for the duration of the optimization.

million. Volume is constrained to a minimum value of 0.2617 cubed reference units, while the lift and pitching moment are constrained to values of $C_L S = 1.7035$ and $C_M \geq -0.17$. Koo[34] had earlier performed a brief examination of the multimodality of this problem when only twist and sectional degrees of freedom are active and concluded that the problem is unimodal.

The purpose of this study is slightly different than that of the previous studies in this work, in that it seeks specifically to investigate the effect that varying the degrees of freedom in a problem has on the presence of multimodality. This is accomplished by beginning with what is believed to be a unimodal problem and gradually increasing the degrees of freedom to see if, or when, multimodality appears.

For this reason the structure of this case differs quite a bit from the previous cases as well: the case is in fact made up of three sub-problems each with progressively more degrees of freedom. The first case solely allows sectional control, then twist is added, and finally taper. In each case a large sample of initial geometries are generated using the sampling algorithm and the parameter values in Table 8 (parameters corresponding to inactive degrees of freedom are neglected). Then 16 viable geometries plus the baseline are optimized using linear and non-linear constraints corresponding as closely as possible to those used in the previous works studying this case. If more than one local optimum is located, the results are noted and the case is finished. If only a single global optimum is found, then an additional 16 viable geometries are attempted. If these also fail to locate any additional optima then the case is concluded to be unimodal for the purposes of this study; if any additional optima are found they are duly noted and categorized. As initial geometries are attempted until 16 or 32 viable samples are located, the number of geometries attempted is generally higher than the final number of tests performed.

Several points regarding the definitions of each case should be clarified. The linear feasibility constraints

Table 7: Blended wing-body estimated grid converged performance values

Value	C_L	C_D	C_L/C_D	% change from optimized baseline	Number of examples
Baseline	0.1485	1.297×10^{-2}	11.456	-	-
Optimized Baseline (1)	**0.3074**	$\mathbf{1.234 \times 10^{-2}}$	**24.907**	**0.00%**	48
Family 2	0.2859	1.471×10^{-2}	19.429	-21.99%	4
Family 3	0.2893	1.277×10^{-2}	22.655	-9.04%	10
Family 4	0.2969	1.328×10^{-2}	22.355	-10.24%	6

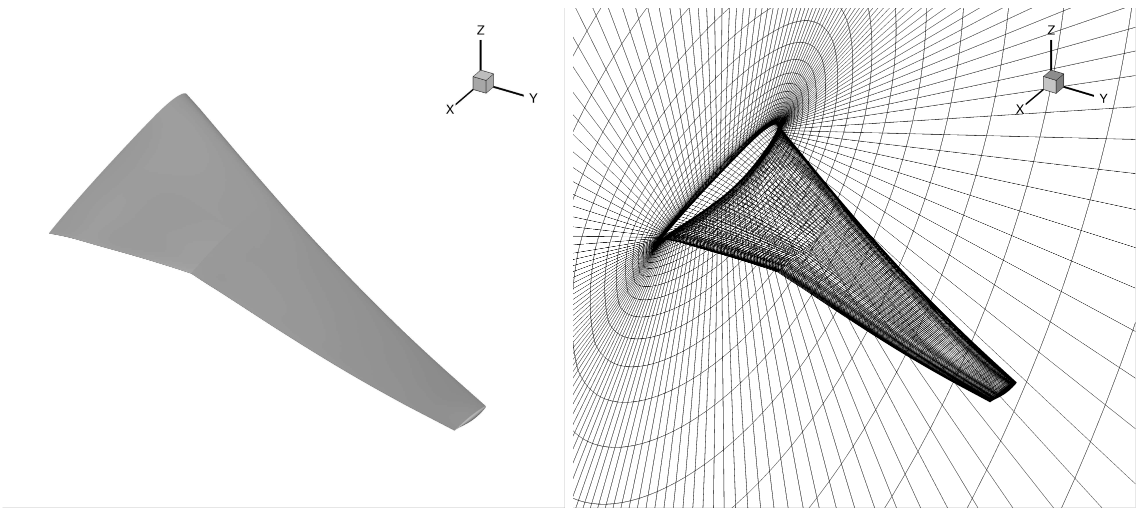

(a) Baseline CRM geometry

(b) Baseline CRM mesh

Figure 12: Baseline geometry for CRM optimization case

developed in this work were applied during sampling to ensure quality initial geometries. However, to better parallel the problem examined by Lee et al., they were not used during optimization, allowing the optimizer significantly more freedom than in the previous problems examined here. One additional nonlinear constraint was used, specifically for the taper case a projected area constraint was enforced, fixing the projected area to the initial value.

Though the specifics of the problem definition of course change from case to case, it can be written generally as:

$$\begin{aligned} \text{minimize} \quad & C_D \\ \text{w.r.t.} \quad & v \\ \text{subject to} \quad & C_L S = 1.7035 \\ & S = 3.407 \\ & V \geq 0.2617 \\ & C_M \geq -0.17 \end{aligned} \tag{15}$$

where C_M is the pitching moment coefficient. For clarity, it should be restated that the projected area constraint is only enforced in the taper case.

Table 8: Linear constraint parameter values for CRM multimodality study

Parameter	Value
B_{low}	$(-0.0545, 0.45, 0.50)$
B_{upp}	$(0.0545, 1.55, 1.50)$
r_{low}	$(-0.0175, -0.1, 0.03, -0.0.03)$
r_{upp}	$(0.0175, 0.1, 0.35, 0.03)$
α_{low}	-3.0^{o}
α_{upp}	5.0^{o}

Table 9: Common Research Model mesh parameters

Grid	Nodes	Off-wall Spacing (mean aerodynamic chord)	Average y^+
L0	925,888	1.5×10^{-6}	0.33

D. Optimization Grid

Due to the lack of interest in the specific performance of each local optimum, the Richardson extrapolation procedure is not carried out for this case and only a single mesh is used. This O-O mesh is shown in Figure 12(b), with the key parameters as in Table 9.

E. Geometry Control

Geometry control is provided by two connected FFD control volumes. Local control is provided by cross sections, five on the first control volume and 11 on the second. Large scale deformations are nominally driven by a pair of axial curves, one connected to each volume. These are in turn controlled by respectively 10 and 19 axial control points; however, for the purposes of this study the axial control is fixed and only cross sectional control is examined. The control system remains constant throughout all tests, though the specific degrees of freedom available are permitted to vary.

F. Multimodality Study Results

The test results tabulated in Table 10 are identified using a three digit code, of the format xxx where each value represents one of the three cross sectional degrees of freedom and takes a zero or one, to denote whether the corresponding degree of freedom is active. From left to right, these values correspond to twist, taper, and section.

In Table 10 the "dominance" column indicates what percentage of the tests converged to the most common local optimum. This provides a quick metric to judge the degree of multimodality fairly across cases that comprise differing quantities of attempted tests, with a lower dominance score suggesting a higher degree of multimodality. This number is not 100% in any case; this is explained by the presence of failed tests (6% for 001, 9% for 101, and 30% for 111), as well as other local optima.

The 001 case permitting only section control and disallowing changes to angle of attack is illustrated in Figures 13, 14, and 15; the primary, and global, optimum, identified by 76% of tests, is depicted in Figure 13, while Figures 14 and 15 depict Optimum 2 and Optimum 3, identified by 15% and 3% of tests, respectively. Each is highly similar to the global optimum, though offering inferior performance, and is differentiated from the global optimum by a single notable geometric feature.

For optimum 1, the most notable change from the baseline geometry can be seen at the root, which is shown superimposed with the baseline root in Figure 13(b), displaying a clear thickening of the cross section. Optimum 2 is primarily differentiated from the global optimum by its root cross section. As can be seen in Figure 14(b), it is characterized by a slightly hooked leading edge relative to the global optimum. The small geometric and performance differences, listed in Table 11, between this and the global optimum, suggest that this may simply be a numerical artifact of the optimization process, rather than a true discrete local

Table 10: Common research model wing multimodality results. Each test is identified using a three digit code of the format xxx. From left to right each digit indicates whether twist, taper, and section freedom are permitted.

Test Code	Tests Performed	Optima Located	Dominance
001	33	3	76%
101	33	3	82%
111	17	6	29%

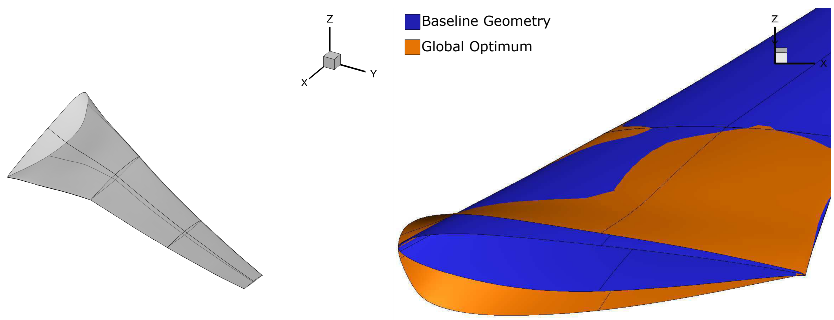

(a) Isometric view (b) Root view versus baseline geometry

Figure 13: Optimum 1 for 001 case

optimum. However, this feature was remarkably resilient to continued optimization, and the design was independently located by several tests, which together are strong suggestions that this design represents a local optimum that has captured the optimizer. Similarly, optimum 3 is highly comparable to optimum 1, though characterized by a small upturned hooking at the wing crank, shown in Figure 15(b). Once again this design did not change appreciably with continued optimization, but as in optimum 2 the relatively small geometric differences and rarity of this result compared to the global optimum makes a definitive categorization of this design difficult. To be conservative, this problem should be considered multimodal, though the design space is highly dominated by the identified global optimum.

In the 101 case, which includes twist and section design freedom - as well as an angle of attack design variable - a similar pattern is seen. The dominant optimum, shown in Figure 16, is located by 82% of the performed tests. Clear differences from the baseline geometry both in twist along the span and section shape throughout the geometry - most noticeably at the root - are apparent. As in the 001 test, two other designs are located by a small number of the performed tests. Optimum 2 (occurrence rate of 6%) is depicted in Figure 17, and Optimum 3 (occurrence rate of 3%) is depicted in Figure 18. Unlike the previous case the geometric differences between these geometries are quite clear and distinctive. The variance in performance (Table 12) is also larger in this case, marking them more definitively as independent local optima. This contradicts the results found by previous authors who have examined this problem. However, 33 tests were attempted before three successful tests were found that did not converge to the global optimum, so it is unsurprising that these other optima might be missed.

Finally, in the 111 case, which includes angle of attack, twist, taper, and section design variables, we find clear evidence of significant multimodality. The most common of the six identified local optima comprises 29% of the results, and all six local optima were found within the first 17 tests. These are shown in Figures 19 to 24. Though all exhibit similar characteristics - particularly a trend towards a BWB design - they are nevertheless clearly differentiated within the bounds of the permitted degrees of freedom, and display a fairly large variation in performance, as can be seen in Table 13. We can conclude that within the scope of the examined cases, section and twist control offer relatively narrow potential for multimodality, while taper

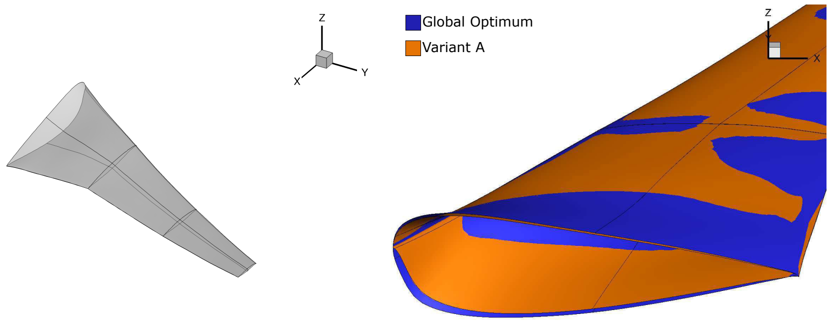

(a) Isometric view (b) Root view versus global optimum

Figure 14: Optimum 2 for 001 case

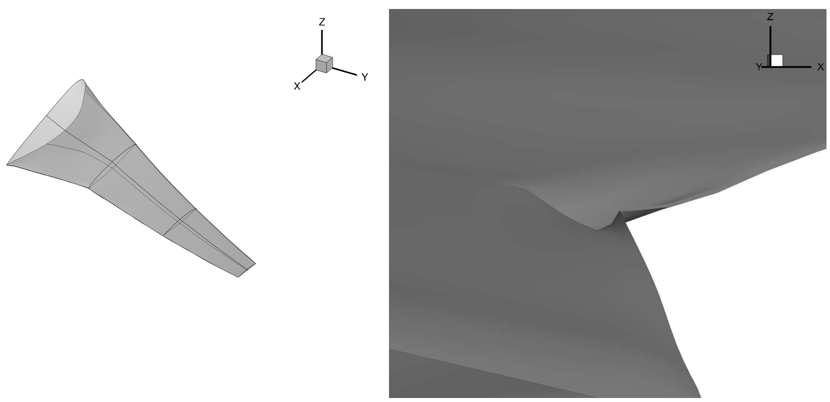

(a) Isometric view (b) Closeup of crank

Figure 15: Optimum 3 for 001 case

adds significant potential. We also conclude that the 101 (twist and section) case is in fact multimodal, in contrast to results by other authors, but this multimodality is of a low degree and as the best optimum is the easiest to locate the odds of an optimizer becoming trapped by a local optimum are quite low.

Table 11: 001 case optimization grid performance values

Value	C_L	C_D	C_L/C_D	% change from optimized baseline	Number of examples
Baseline	0.500	2.21×10^{-2}	22.61	-	-
Optimized Baseline (1)	**0.500**	$\mathbf{1.98 \times 10^{-2}}$	**25.25**	**0.00%**	25
Optimum 2	0.500	1.99×10^{-2}	25.08	-0.68%	5
Optimum 3	0.500	2.00×10^{-2}	25.04	-0.86%	1

Table 12: 101 case optimization grid performance values

Value	C_L	C_D	C_L/C_D	% change from optimized baseline	Number of examples
Baseline	0.500	2.21×10^{-2}	22.61	-	-
Optimized Baseline (1)	**0.500**	$\mathbf{1.99 \times 10^{-2}}$	**25.15**	**0.00%**	27
Optimum 2	0.500	2.03×10^{-2}	24.66	-1.97%	2
Optimum 3	0.500	2.00×10^{-2}	25.02	-0.51%	1

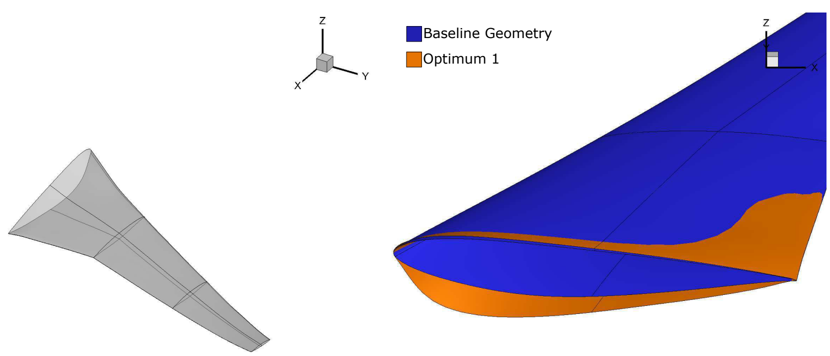

(a) Isometric view (b) Root view versus baseline

Figure 16: Optimum 1 for 101 case

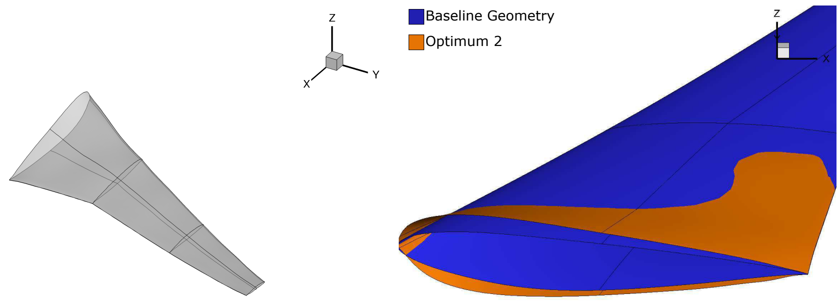

(a) Isometric view (b) Root view versus baseline

Figure 17: Optimum 2 for 101 case

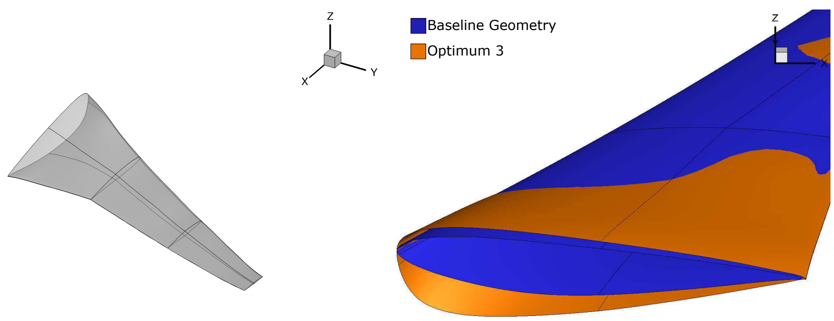

(a) Isometric view (b) Root view versus baseline

Figure 18: Optimum 3 for 101 case

Table 13: 111 case optimization grid performance values

Value	C_L	C_D	C_L/C_D	% change from optimized baseline	Number of examples
Baseline	0.500	2.21×10^{-2}	22.61	-	-
Optimized Baseline (1)	0.500	1.82×10^{-2}	27.54	0.00%	5
Optimum 2	0.500	1.81×10^{-2}	27.56	0.06%	3
Optimum 3	0.500	1.84×10^{-2}	27.24	−1.07%	1
Optimum 4	0.500	1.82×10^{-2}	27.40	−0.50%	1
Optimum 5	0.500	1.82×10^{-2}	27.47	−0.23%	1
Optimum 6	**0.500**	$\mathbf{1.81 \times 10^{-2}}$	**27.62**	**0.28%**	**1**

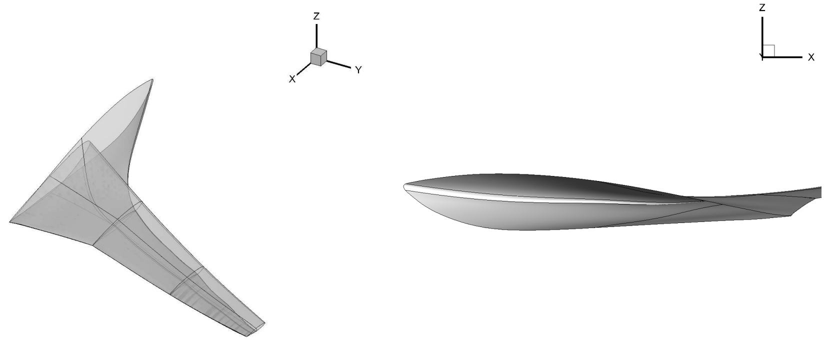

(a) Isometric view versus baseline (b) Root view

Figure 19: Optimum 1 for 111 case

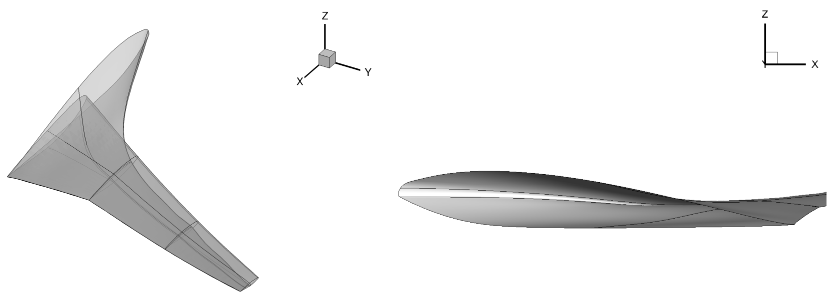

(a) Isometric view versus baseline (b) Root view

Figure 20: Optimum 2 for 111 case

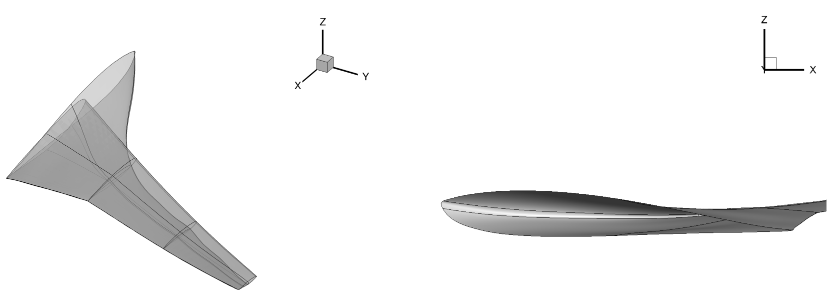

(a) Isometric view versus baseline (b) Root view

Figure 21: Optimum 3 for 111 case

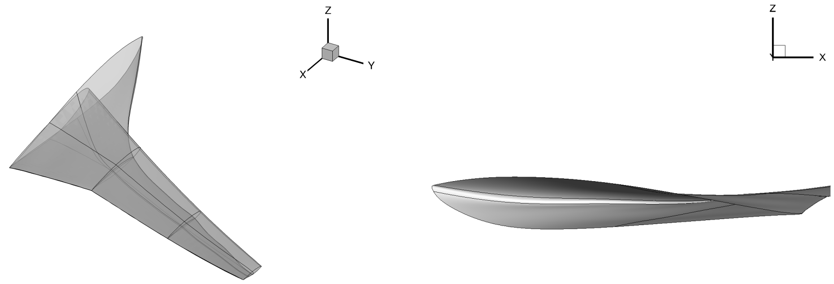

(a) Isometric view versus baseline (b) Root view

Figure 22: Optimum 4 for 111 case

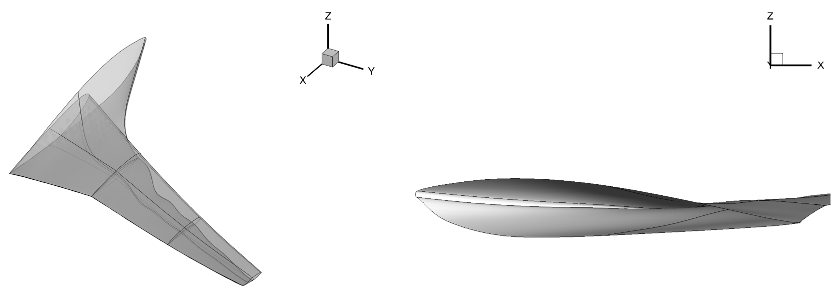

(a) Isometric view versus baseline (b) Root view

Figure 23: Optimum 5 for 111 case

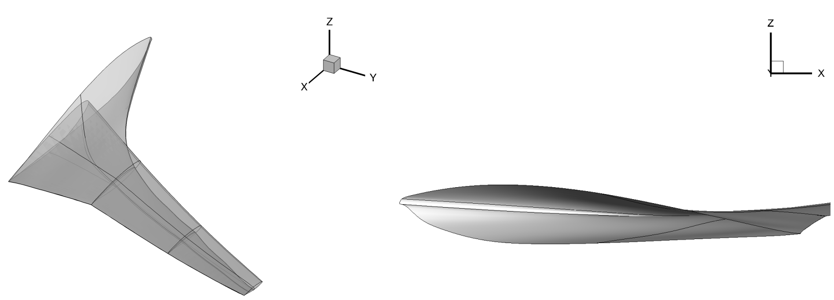

(a) Isometric view versus baseline (b) Root view

Figure 24: Optimum 6 for 111 case

V. Conclusion

An improvement has been presented to an efficient multi-start gradient-based algorithm for multimodal aerodynamic shape optimization problems. The use of FFD-based geometry control simplifies the development of the cascading constraints used in generating the initial samples and constraining the subsequent optimization. Tests investigated the impact of various degrees of freedom on the degree of multimodality present. Multimodality was demonstrated in the design spaces for viscous exploratory optimizations of wings in both transonic and subsonic conditions, while a more constrained transonic BWB case was tentatively found to be unimodal. The multimodality study utilizing the CRM wing in transonic viscous flows suggests that taper holds the strongest potential for multimodality out of the examined degrees of freedom. It was also concluded that, in contrast to previously published results, the common twist and section CRM case is in fact multimodal; however, the design space is overwhelmingly dominated by the global optimum making it unlikely that the optimizer would be captured by one of the outlying optima.

VI. Acknowledgments

The authors would like to gratefully acknowledge the financial support of the University of Toronto, as well the Ontario Graduate Scholarship program. Computations were performed on the General Purpose Cluster supercomputer at the SciNet HPC Consortium. SciNet is funded by: the Canada Foundation for Innovation under the auspices of Compute Canada; the Government of Ontario; Ontario Research fund - Research Excellence; and the University of Toronto.

References

[1]IATA, "IATA Annual Review 2015," 2015.

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

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

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

[5]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.

[6]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.

[7]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.

[8]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.

[9]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.

[10]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.

[11]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.

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

[13]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.

[14]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.

[15]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.

[16]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.

[17]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.

[18]Gagnon, H. and Zingg, D. W., "Two-level free-form deformation for high-fidelity aerodynamic shape optimization," *12th AIAA Aviation Technology, Integration and Operations (ATIO) Conference and 14th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, 2012.

[19]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.

[20]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.

[21]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.

[22]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.

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

[24]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.

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

[26]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.

[27]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.

[28]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.

[29]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., et al., "Summary of data from the fifth AIAA CFD drag prediction workshop," 2013.

[30]Koo, D. and Zingg, D. W., "Progress in Aerodynamic Shape Optimization Based on the Reynolds-Averaged Navier-Stokes Equations," *54th AIAA Aerospace Sciences Meeting*, No. 2016-1292, 2016.

[31]Roache, P. J., *Verification and validation in computational science and engineering*, Hermosa, Albuquerque, NM, 1998.

[32]Lee, C., Koo, D., and Zingg, D. W., "Comparison of B-Spline Surface and Fee-form Deformation Geometry Control for Aerodynamic Optimization," *AIAA Journal*, Vol. 55, No. 1, pp. 228–240.

[33]Reist, T. A. and Zingg, D. W., "Optimization of the Aerodynamic Performance of Regional and Wide-Body-Class Blended Wing-Body Aircraft," *33rd AIAA Applied Aerodynamics Conference*, No. 2015-3292, 2015, To appear in Journal of Aircraft, 2017.

[34]Koo, D., "An Investigation into Aerodynamic Shape Optimization of Planar and Nonplanar Wings," *AIAA Journal*, Submitted.