

I. Introduction

One of the primary challenges of hypersonic airbreathing vehicles is the design of the engine inlet. In order for the overall propulsion system to be effective, the inlet must compress the freestream airflow to the desired Mach number while minimizing losses. For a vehicle that accelerates across a range of Mach numbers, single point optimization at the expense of other conditions is insufficient. Rather, the design should be optimized for the entire trajectory. Additionally, performance margins are typically small enough that uncertainty must be accounted for and preferably minimized. The first step in reducing uncertainty at the vehicle level is reducing uncertainty at the component level. Component level uncertainty can be reduced by employing robust optimization, or design under uncertainty, instead of classical deterministic optimization.

The first attempt at optimizing supersonic inlet designs was made by Oswatitsch [1]. Oswatitsch analytically optimized the design of a ramjet inlet through the use of Lagrange multipliers and inviscid gas dynamic relations and obtained the result that an optimally designed inlet should have successively increasing ramp angles as the Mach number is reduced such that shock strength remains constant. This design strategy minimizes the losses incurred by excessively large shocks. After Oswatitsch, Smart [2] extended the analysis from ramjets to scramjets by once again using Lagrange multipliers and inviscid gas dynamic relations. Hasegawa and Knight [3] used computational fluid dynamics (CFD) within a gradient-based optimization loop to maximize total pressure recovery. Brown et al. [4] performed robust design with an evolutionary algorithm and inviscid analysis, then verified their results with CFD. However, they wrote that "the use of a numerically intensive CFD code embedded in the optimization procedure is not feasible here" [4].

The purpose of the present work is to apply the techniques of optimization under uncertainty to the design of scramjet inlets. Multiple fidelity levels are used in the analysis. Low-fidelity analysis consists of analytical oblique shock relations, while high-fidelity analysis consists of viscous CFD. The low-fidelity analysis is able to capture most of the physics of the problem in a time-efficient manner, but there are some effects that are only captured with high-fidelity analysis. Multifidelity uncertainty quantification is used to reduce the computational expense of uncertainty quantification. Low-fidelity analysis, grid generation, high-fidelity CFD, and uncertainty quantification are all performed entirely within the optimization loop. The remainder of the paper is organized as follows. Section II describes the theory and methodology used in this work, including the low-fidelity analysis, the high-fidelity analysis, the uncertainty quantification technique, the optimization algorithm, and a numerical uncertainty estimation technique. Section III describes the computational setup and formulation of the optimization problem, and Section IV presents the results of several studies. Finally, Section V describes the conclusions of this work.

II. Methodology

The following sections describe the background theory used in the present work. Section II.A provides the compressible gas dynamic relations which form the foundation of the low-fidelity analysis technique, while Section II.B describes the details of the computational fluid dynamics solver used as a high-fidelity analysis technique. Section II.C describes multifidelity uncertainty quantification through the use of non-intrusive point collocation polynomial chaos expansions. Section II.D describes the optimization algorithm used in the present work, Efficient Global Optimization (EGO). Finally, Section II.E describes a numerical uncertainty estimation technique used for interpreting the EGO results.

A. Compressible Gas Dynamics

For a two-dimensional scramjet inlet composed of angled ramps, the key inviscid physics are given by the well-known oblique shock relations for a calorically perfect gas [5]. Given an upstream Mach number M_1 and a ramp angle θ , the oblique shock wave angle β can be calculated by iterating on the well-known θ - β - M function,

$$\tan(\theta) = 2 \cot(\beta) \left[\frac{M_1^2 \sin^2(\beta) - 1}{M_1^2 (\gamma + 2 \cos(2\beta)) + 2} \right]. \quad (1)$$

There is also an equivalent formulation, known as the β - θ - M relation, which allows for explicit numerical calculation of β without iterating. The β - θ - M relation is

$$\tan(\beta) = \frac{M_1^2 - 1 + 2\lambda \cos[(4\pi\delta + \cos^{-1}(\chi))/3]}{3 \left(1 + \frac{\gamma-1}{2} M_1^2 \right) \tan(\theta)}, \quad (2)$$

where the weak shock solution is given by $\delta = 1$, the strong shock solution is given by $\delta = 0$, and

$$\lambda = \left[(M_1^2 - 1)^2 - 3 \left(1 + \frac{\gamma - 1}{2} M_1^2 \right) \left(1 + \frac{\gamma + 1}{2} M_1^2 \right) \tan^2(\theta) \right]^{1/2}, \quad (3)$$

$$\chi = \frac{(M_1^2 - 1)^3 - 9 \left(1 + \frac{\gamma - 1}{2} M_1^2 \right) \left(1 + \frac{\gamma - 1}{2} M_1^2 + \frac{\gamma + 1}{4} M_1^4 \right) \tan^2(\theta)}{\lambda^3}. \quad (4)$$

With the shock wave angle β known from either of the above approaches, the Mach number normal to the shock, $M_{n,1}$, can be found with

$$M_{n,1} = M_1 \sin(\beta). \quad (5)$$

The properties behind the shock, expressed as ratios of the pre-shock values, are

$$\frac{p_2}{p_1} = 1 + \frac{2\gamma}{\gamma + 1} (M_{n,1}^2 - 1), \quad (6)$$

$$\frac{\rho_2}{\rho_1} = \frac{(\gamma + 1) M_{n,1}^2}{2 + (\gamma - 1) M_{n,1}^2}, \quad (7)$$

and

$$\frac{T_2}{T_1} = \left(\frac{p_2}{p_1} \right) \left(\frac{\rho_1}{\rho_2} \right). \quad (8)$$

The Mach number behind and normal to the shock, $M_{n,2}$, is

$$M_{n,2} = \left(\frac{\rho_1}{\rho_2} \right) \left(\frac{T_1}{T_2} \right)^{1/2}, \quad (9)$$

and the Mach number behind the shock is

$$M_2 = \frac{M_{n,2}}{\sin(\beta - \theta)}. \quad (10)$$

Using the Mach number and static property ratios behind the shock, the total properties behind the shock are calculated with the ratios

$$\frac{T_t}{T} = 1 + \frac{\gamma - 1}{2} M^2, \quad (11)$$

$$\frac{p_t}{p} = \left(1 + \frac{\gamma - 1}{2} M^2 \right)^{\frac{\gamma}{\gamma - 1}}, \quad (12)$$

and

$$\frac{\rho_t}{\rho} = \left(1 + \frac{\gamma - 1}{2} M^2 \right)^{\frac{1}{\gamma - 1}}. \quad (13)$$

Equations (1) through (13) can be applied iteratively for each ramp in the inlet to compute the throat conditions, assuming that there are no expansion waves or reflected shocks, and that the gas is calorically perfect. This approach does not account for viscous losses, but the inviscid approximation of inlet exit conditions is acceptable, especially considering its execution speed compared to viscous CFD.

B. Computational Fluid Dynamics Solver

The CFD software used in this study is Viscous Upwind Algorithm for Complex Flow Analysis (VULCAN-CFD) [6], which is a multiblock, cell-centered, finite-volume solver widely used for high-speed flow simulations. In the present work, viscous simulations are performed using the Reynolds-averaged Navier-Stokes (RANS) equations [7] for thermally perfect air. The Low-Diffusion Flux-Split Scheme (LDFSS) of Edwards [8] is used, and interpolation is done using the Monotone Upstream-Centered Scheme for Conservation Laws (MUSCL) [9]. The van Leer flux limiter [9] is used to prevent oscillations in regions near shocks. Turbulence is modeled using the Menter-BSL model [10]

with wall functions enabled. The implicit incomplete lower-upper (ILU) time-stepping scheme [11] is used with a Courant-Friedrichs-Lewy (CFL) number of 50 to solve the governing equations.


The two-dimensional grids used in this study have approximately 120,000 cells and are designed to have a y^+ value at the walls of roughly 200, which is within the log-law portion of the boundary layer and sufficient for using wall functions. Inflow conditions are set as Mach 6 freestream air with a dynamic pressure of 1500 pounds per square foot (psf) and an angle of attack of 0 degrees. The boundary conditions at all surfaces are no-slip walls with a fixed isothermal temperature of 750 Kelvin. The grids are split into multiple blocks to allow for parallel processing. Grids are generated automatically within the optimization loop using Pointwise [12] and custom scripts.

Because CFD executes within an optimization loop, without the opportunity for manual inspection, a conservative approach to ensuring iterative convergence is necessary. A sample case using an optimized design from the low-fidelity analysis of Section II.A was run until the L_2 -norm of the residuals, the percent difference between mass inflow and outflow, and the integrated normal and shear force (F_x and S_x) at the wall boundaries had all converged. Plots of these quantities from the sample case are shown in Figure 1. Convergence occurred at roughly 6,000 iterations. To ensure adequate convergence for all optimization cases, the CFD simulations within the optimization loop are set to run for 10,000 iterations using the same settings.

Once a CFD simulation converges, post-processing is performed. A one-dimensionalization algorithm, Massflow-3D, is run to compute a mass-weighted average of all flow properties as a function of the streamwise coordinate [13]. The location of the inlet throat is known, and the one-dimensionalized quantities of interest are extracted for the throat plane. Scalar values produced from the one-dimensionalization are input to the optimization objective function. Unlike the low-fidelity analysis of Section II.A, this high-fidelity method can account for the effects of viscosity, varying specific heats, turbulence, boundary layer separation, and the presence of reflected shocks and expansion waves.

C. Uncertainty Quantification via Polynomial Chaos Expansion

Uncertainty quantification (UQ) in computational modeling is the process of determining how uncertainty in input quantities, numerical errors, and modeling assumptions and limitations will affect the output responses [14]. Uncertainty can be classified into two types: aleatory and epistemic [15]. Aleatory uncertainty, also called irreducible uncertainty, is inherent to a system and cannot be reduced. Examples of aleatory uncertainty include sensor noise and fluctuations in wind speed. Epistemic uncertainty, sometimes called reducible uncertainty, results from either a lack of knowledge or an intentional simplification of modeling. An example of epistemic uncertainty is using a RANS turbulence model, rather than resolving all of the scales of turbulence as in direct numerical simulation (DNS). In principle, epistemic uncertainty can be reduced, however, doing so is not always feasible or practical due to lack of budget, computational time, experimental data, or other finite resources. Aleatory uncertainty is well-handled by probabilistic methods, while



(a) L_2 -norm of the residuals and mass flow error.

(b) Integrated normal and shear forces at the walls.

Fig. 1 CFD iterative convergence for a representative sample case.

epistemic uncertainty is less straightforward to model.

The technique used to forward propagate input uncertainties in the present work is the polynomial chaos expansion (PCE). Polynomial chaos expansions start with the Cameron-Martin theorem [16] by representing a finite variance random variable or response function R as an infinite expansion [17]

$$R = a_0 B_0 + \sum_{i_1=1}^{\infty} a_{i_1} B_1(\xi_{i_1}) + \sum_{i_1=1}^{\infty} \sum_{i_2=1}^{i_1} a_{i_1 i_2} B_2(\xi_{i_1}, \xi_{i_2}) + \sum_{i_1=1}^{\infty} \sum_{i_2=1}^{i_1} \sum_{i_3=1}^{i_2} a_{i_1 i_2 i_3} B_3(\xi_{i_1}, \xi_{i_2}, \xi_{i_3}) + \dots \quad (14)$$

with real coefficients a and interaction polynomial terms B . As stated by Eldred [14], the above expression with polynomial order-based indexing can be replaced with term-based indexing to yield a simpler expression,

$$R(\mathbf{D}, \boldsymbol{\xi}) = \sum_{j=0}^{\infty} \alpha_j(\mathbf{D}) \Psi_j(\boldsymbol{\xi}), \quad (15)$$

where α_j is a deterministic component as a function of deterministic variables \mathbf{D} , and Ψ_j is a random variable basis function using the random variables $\boldsymbol{\xi}$. With polynomial chaos, the basis functions of the random variables are determined by using the Askey key [18], which relates the distributions of the uncertain parameters to specific orthogonal polynomials. In order to make this approach feasible, the above infinite series is truncated at a finite number of dominant terms,

$$R(\mathbf{D}, \boldsymbol{\xi}) \approx P_{q,n} = \sum_{j=0}^{N_t-1} \alpha_j(\mathbf{D}) \Psi_j(\boldsymbol{\xi}) \quad (16)$$

where $P_{q,n}$ is the PCE of highest order q for n random dimensions.

With the random variable basis functions Ψ_j known from the Askey key, the expansion coefficients α_j must now be determined. This is done through non-intrusive point collocation [19, 20], which requires only the system response values at P dispersed sample points. If the sample points are linearly independent, then the result is the following linear system

$$\begin{bmatrix} R(\mathbf{D}, \boldsymbol{\xi}_0) \\ R(\mathbf{D}, \boldsymbol{\xi}_1) \\ \vdots \\ R(\mathbf{D}, \boldsymbol{\xi}_P) \end{bmatrix} \approx \begin{bmatrix} \Psi_0(\boldsymbol{\xi}_0) & \Psi_1(\boldsymbol{\xi}_0) & \cdots & \Psi_P(\boldsymbol{\xi}_0) \\ \Psi_0(\boldsymbol{\xi}_1) & \Psi_1(\boldsymbol{\xi}_1) & \cdots & \Psi_P(\boldsymbol{\xi}_1) \\ \vdots & \vdots & \ddots & \vdots \\ \Psi_0(\boldsymbol{\xi}_P) & \Psi_1(\boldsymbol{\xi}_P) & \cdots & \Psi_P(\boldsymbol{\xi}_P) \end{bmatrix} \begin{bmatrix} \alpha_0 \\ \alpha_1 \\ \vdots \\ \alpha_P \end{bmatrix}. \quad (17)$$

Note that the minimum value for P is $P = N_t - 1$ where the number of terms, N_t , is given by

$$N_t = OSR \frac{(n+q)!}{n!q!}, \quad (18)$$

where an oversampling ratio $OSR \geq 1.0$ may be used. If more sample points are available through using an oversampling ratio greater than 1, then the system can be solved with a least squares method. In this work, an oversampling ratio of 1.5 is used for the purpose of providing additional data. For this work, including up to second-order polynomial terms, $q = 2$, was found to be sufficient based on cross-validation and test points.

Once the α_j coefficients have been computed, the response mean and variance can be determined from closed form analytical expressions [21]

$$\mu = \langle R \rangle \approx \sum_{j=0}^P \alpha_j \langle \Psi_j(\boldsymbol{\xi}) \rangle = \alpha_0 \quad (19)$$

$$\sigma^2 = \langle (R - \mu)^2 \rangle \approx \langle \left(\sum_{j=1}^P \alpha_j \Psi_j(\boldsymbol{\xi}) \right)^2 \rangle = \sum_{j=1}^P \sum_{k=1}^P \alpha_j \alpha_k \langle \Psi_j(\boldsymbol{\xi}) \Psi_k(\boldsymbol{\xi}) \rangle = \sum_{j=1}^P \alpha_j^2 \langle \Psi_j(\boldsymbol{\xi})^2 \rangle \quad (20)$$

If multiple analysis fidelity levels are available, then multifidelity UQ can be performed. The advantage of multifidelity UQ is in achieving results that closely approximate those of purely high-fidelity UQ at a substantially reduced computational cost. Reducing the cost of UQ makes multifidelity analysis appealing for use within robust

optimization due to the large computational burden. As stated by West and Philips [20], the goal of multifidelity PCE is to use a small number of high-fidelity training points to correct a low-fidelity PCE model. If the high-fidelity correction is additive, then

$$R_H = R_L + C \approx P_{q,n} + C \quad (21)$$

where the high-fidelity response R_H is equal to the low-fidelity response R_L plus a correction term C . Using the PCE approach described previously, the low-fidelity response can be approximated by a polynomial chaos expansion of order q and random dimensions n . Rearranging Eq. (21), the additive correction can be expressed as

$$C = R_H - R_L \approx P_{r,n} \quad (22)$$

where $P_{r,n}$ is a polynomial chaos expansion of order r , with $r < q$.

Observe that the above discussion mentioned that a PCE model of order $q = 2$ was found to be sufficient for capturing the high-fidelity physics of this problem, thus using a PCE of order $r < q$ in the multifidelity approach may at first seem concerning. However, the PCE of Eq. (22) does not need to capture all of the physics of the high-fidelity model. The PCE model of order q , built from the low-fidelity model, should already capture most of the physics. The PCE model of order r thus only needs to capture the difference in physics between the fidelity levels in order to approximate R_H ,

$$R_H = R_L + C \approx P_{q,n} + C \approx P_{q,n} + P_{r,n}. \quad (23)$$

With this approximation of R_H , a system of equations similar to Eq. (17) can be constructed and the response statistics solved for using Eqs. (19) and (20). Note that the goal of multifidelity UQ is to reduce the number of high-fidelity model evaluations, computed from Eq. (18). Therefore, if $r \geq q$, no high-fidelity evaluations are saved compared to purely high-fidelity UQ, and the multifidelity approach does not contribute any value. Further discussion of polynomial chaos expansions and multifidelity techniques are available from Ng and Eldred [19], Ng and Willcox [22], and Peherstorfer et al [23]. An extension of multifidelity techniques for when high and low fidelity models have different uncertainties is available from West and Phillips [20].

D. Optimization Algorithm

Classical gradient-based optimization works well when the function being optimized is convex and gradients are efficiently and accurately obtained. However, many practical problems can be non-convex and multimodal such that local minima exist and cause difficulty in finding the global optimum. When a gradient-based optimizer is initialized with a poor starting point, it may descend into a local minima and terminate, therefore missing the global optimum. Another issue with gradient-based algorithms, specific to robust design, is that the computation of gradients of uncertainty statistics can become prohibitively expensive. The large cost is caused by the need to propagate uncertainty for each finite difference applied to the design variables. Table 1 below demonstrates the large cost by carrying out calculations for the number of function evaluations needed for various optimization methods. Robust gradient-based optimization with second-order accurate gradients can easily reach hundreds of function evaluations per iteration, and the optimization algorithm itself may take between dozens and hundreds of iterations to converge. Expensive function evaluations, such as CFD or finite element analysis (FEA), exacerbate the problem. Both of these problems are jointly addressed by using

Table 1 Number of function evaluations required at each optimization iteration.

Algorithm Type	Design variations per iteration	Function evaluations per variation	Total function evaluations per iteration	Evaluations per iteration for sample problem, $N_d = 7$ and $N_t = 15$
Deterministic gradient-based, $O(\Delta\mathbf{x})$	$(1 + N_d)$	1	$(1 + N_d)$	8
Deterministic gradient-based, $O(\Delta\mathbf{x}^2)$	$(1 + 2N_d)$	1	$(1 + 2N_d)$	15
Deterministic derivative-free	1	1	1	1
Robust gradient-based, $O(\Delta\mathbf{x})$	$(1 + N_d)$	$(1 + N_t)$	$(1 + N_d)(1 + N_t)$	128
Robust gradient-based, $O(\Delta\mathbf{x}^2)$	$(1 + 2N_d)$	$(1 + N_t)$	$(1 + 2N_d)(1 + N_t)$	240
Robust derivative-free	1	$(1 + N_t)$	$(1 + N_t)$	16

a global surrogate-based derivative-free optimization algorithm, which searches the entire design space for a global optimum without the need to compute gradients. In this work, the Efficient Global Optimization [24, 25], or EGO, algorithm is used. The specific implementation of EGO used is that included within the Dakota software package [26]. The remainder of this section provides a brief description of the EGO algorithm.

An important requirement on global optimization methods is the need to balance between improving the current best-known value, termed exploitation, with sampling points of the design space that are not well-known, termed exploration. EGO accomplishes this balance with an expected improvement function (EIF). The EIF is the expectation that a specific point in the feasible design space will provide a better value than the current known best value. The predicted value for an unknown point comes from a Gaussian process (GP) model, which provides a Gaussian error distribution for the point being predicted. This distribution permits the expectation to be computed. Points with near-optimal predicted values and low variances have a good chance of yielding an improvement over the current best known value, as do points with poor predicted values and large variances. Sampling the former point allows EGO to exploit current information known about the function, while sampling the latter point allows EGO to explore new territory in the design space. Thus, by successively sampling whichever point in the design space currently has the maximum EIF, the EGO algorithm gradually builds an improving surrogate GP model of the function. The algorithm continues to iterate until the maximum expected improvement is less than 1% of the current best known value.

Note that the original implementation of EGO in Jones et al. [24] uses a branch-and-bound method to maximize the EIF. The implementation of EGO in Dakota uses the DIRECT algorithm [27] instead. For further information on EGO, refer to the original work by Jones et al [24].

E. Optimizer Early Exit Error and Uncertainty Estimation

The use of a global surrogate-based derivative-free optimizer has many advantages compared to gradient-based optimization. However, there is also a disadvantage that must be accounted for. Gradient-based optimization algorithms, such as steepest descent, conjugate gradient, and sequential quadratic programming (SQP) typically run until a convergence tolerance on the order of 1×10^{-4} to 1×10^{-6} is reached [28]. As described in the previous section, the Dakota implementation of EGO runs until the maximum expected improvement is less than 1% of the current best known value. The user is unable to change the convergence criteria in the Dakota implementation of EGO. As a result of the difference in convergence criteria, EGO will typically find the valley that contains the global minimum, but it does not have the required level of tolerance to descend to the bottom of the valley. The difference between the value returned by EGO and the true optimum at the bottom of the valley is termed early exit error due to it being caused by EGO terminating earlier than a gradient-based optimizer typically would.

Results from an EGO optimization should have epistemic uncertainty bands applied to them to account for this early exit error. These uncertainty bands can be calculated by first quantifying the early exit error ϵ_{ee} . This is done by using the output design point of EGO as the initial design point of a gradient-based optimizer such as SQP. The gradient-based optimizer is run until convergence and returns the design point and functionals at the minimum of the valley found by EGO, considered to be the true optimum. The difference between the EGO results and the gradient-based results is thus the early exit error,

$$\epsilon_{ee} = f_{EGO} - f_{SQP}, \quad (24)$$

where f is the quantity of interest being examined. An estimated uncertainty band U can be calculated for some quantity by taking an appropriate norm of the early exit error, such as the L_2 or L_∞ norm, and then multiplying by a factor of safety F_s

$$U = F_s L_\infty(\epsilon_{ee}). \quad (25)$$

As discussed in Section II.D, running a gradient-based optimizer is not always feasible, especially for robust design. In this situation, the early exit error should be quantified for at least one case using the same objective function and constraints as the desired optimization, and then the results can be applied to other cases. This is similar in practice to the estimation of numerical uncertainty due to discretization error [15, 29]. In the present work, early exit uncertainty is calculated using deterministic optimization with the L_∞ norm, a factor of safety $F_s = 1.5$, and then the epistemic uncertainty band is applied to EGO results.

III. Computational Setup

With the background theory described in the previous section, the following section details the solution process for a specific sample problem. The mathematical formulation of the optimization problem is described in Section III.A,

while the simulation setup and software configuration are detailed in Section III.B.

A. Optimization Problem Formulation

The problem to be solved in the present work is, given some freestream condition, to find the optimal set of 2D planar ramp angles for a scramjet inlet. There are seven inlet ramp angles, divided up into four external compression and three internal compression ramps. This problem is first considered deterministically and then stochastically. The performance objectives are to compress the freestream air to a target Mach number with minimum total pressure losses by the inlet throat. The deterministic problem is formulated mathematically as

$$\min J_{Det} = W_M \left[\frac{(M_{throat} - M_{target})}{(M_{worst} - M_{target})} \right]^2 + W_\Pi \left[\frac{(\Pi_{throat} - \Pi_{target})}{(\Pi_{worst} - \Pi_{target})} \right]^2, \quad (26a)$$

$$\text{subject to } \sum_{external} \theta - \sum_{internal} \theta = 0, \quad (26b)$$

$$2^\circ \leq \theta_i \leq 12^\circ, \quad (26c)$$

where J_{det} is the deterministic problem objective function and the inlet ramp angles, θ_i , are the design variables. Π_{throat} and M_{throat} are the total pressure recovery and the Mach number at the throat plane, respectively, and are the values produced by the subsequent analysis. The *target* subscript refers to the desired quantities for the optimizer to obtain for Π_{throat} and M_{throat} . W_Π and W_M are deterministic weighting coefficients that sum up to 1,

$$W_M + W_\Pi = 1, \quad (27)$$

and are used to give the design engineer control over which term in the objective function is to be prioritized more. The *worst* subscript refers to a feasible value that an objective can achieve that is furthest away from the target value [30]. This worst value is used to normalize the two terms in the objective function to between 0 and 1, ensuring that they are of the same magnitude. Constraint Eq. (26b) ensures that the sum of external inlet angles $\sum_{external} \theta = \theta_1 + \theta_2 + \theta_3 + \theta_4$ is equal to the sum of internal inlet angles $\sum_{internal} \theta = \theta_5 + \theta_6 + \theta_7$ to result in a horizontal inlet exit, parallel to the freestream. Constraint Eq. (26c) applies upper and lower bounds on the design variables.

Target values for the optimizer are $\Pi_{target} = 1.0$ and $M_{target} = 3.0$. These values are chosen because they represent the maximum possible compression efficiency and a moderate amount of compression. The worst values for an inlet to produce are calculated by taking the total pressure recovery that would result from a normal shock at the given freestream Mach number, and a Mach number that is unchanged from its freestream value. For this specific problem, these values are roughly $\Pi_{worst} = 0.03$ and $M_{worst} = 6.0$. An additional parameter is used to provide the physical scale of the problem, x_{cle} , which is the x coordinate of the cowl leading edge as measured from the vehicle leading edge. The sample problem in this work uses a value of $x_{cle} = 700$ inches. A representative image of the type of inlet considered in this work is shown in Figure 2.

For the stochastic problem, uncertainty in the freestream conditions is considered. This allows variance in the trajectory to be accounted. All uncertain parameters are modeled as aleatory. The sample problem of the current work uses Gaussian distributions for freestream Mach number, angle of attack, and dynamic pressure. Their means are the same as the values used in the deterministic problem, with standard deviations as shown in Table 2. Using the UQ techniques described in Section II.C, the mean and standard deviation of M_{throat} and Π_{throat} are calculated for the distributions of Table 2. This yields the values μ_M , σ_M , μ_Π , and σ_Π which are used to calculate the robust objective function J_{Robust} . The design under uncertainty problem is formulated as

$$\min J_{Robust} = W_M \left[W_\mu \left(\frac{(\mu_M - M_{target})}{(M_{worst} - M_{target})} \right)^2 + W_\sigma \sigma_M \right] + W_\Pi \left[W_\mu \left(\frac{(\mu_\Pi - \Pi_{target})}{(\Pi_{worst} - \Pi_{target})} \right)^2 + W_\sigma \sigma_\Pi \right], \quad (28a)$$

$$\text{subject to } \sum_{external} \theta - \sum_{internal} \theta = 0, \quad (28b)$$

$$2^\circ \leq \theta_i \leq 12^\circ, \quad (28c)$$

where W_μ and W_σ are weighting coefficients that allow the design engineer to choose between prioritizing higher mean performance or minimizing variance of performance. Similar to the deterministic weighting coefficients, the robust weighting coefficients are constrained such that

$$W_\mu + W_\sigma = 1. \quad (29)$$

Table 2 Freestream condition uncertainties.

Freestream Parameter	μ	σ
Mach number	6.0	0.1
Angle of attack (deg)	0.0	0.25
Dynamic pressure (psf)	1500.0	15.0

Observe that the robust objective function is of similar form to the deterministic objective function, but the throat quantities have been replaced with their mean values produced from uncertainty propagation. Additionally, there is another tradeoff for the design engineer between high mean performance and low variance of performance. This robust tradeoff is also affected by the deterministic tradeoff, such that the choice of deterministic weighting coefficients can influence whether low variance of M_{throat} or low variance of Π_{throat} is prioritized more.

B. Simulation Setup

The simulation setup for solving the inlet design under uncertainty problem is composed of several key components. The first component is the optimization algorithm, EGO, as implemented in the Dakota software package [26]. The next component is an optimization driver script, which controls fidelity levels, optimization settings, and variable values. The final component is the analysis wrapper, which performs the inlet analysis for a specified geometry at the requested fidelity level. A flowchart of the simulation system is shown in Figure 3.

The first step is for Dakota to set the design variables for the current iteration. The design variables are passed into the optimization driver, which then sets all other parameters in the problem. These parameters include physical values, the distributions of uncertain parameters, the weighting coefficients for both the deterministic and robust objective functions, and software settings such as the fidelity level or whether a deterministic or robust optimization should be performed. Once the optimization driver has set the problem parameters, it calls out to the analysis wrapper. Regardless of the fidelity level selected, the analysis wrapper runs the low-fidelity analysis because the high-fidelity analysis requires geometry data output by the low-fidelity analysis.

The low-fidelity analysis uses the compressible gas dynamic relations of Section II.A to compute the total pressure recovery and the Mach number at the throat of the inlet. This is done using a non-intersecting non-reflecting oblique shock, or NINROS, design strategy [31]. With a fixed x_{cle} value and a known θ_1 of the vehicle leading edge, the oblique shock angle from the leading edge β_1 can be computed. The y coordinate of the cowl leading edge, y_{cle} , is placed such that the oblique shock of β_1 impinges upon it at x_{cle} , referred to as the shock-on-lip condition. Using the relations of Section II.A and the design variable θ_2 , the value of β_2 is calculated. An oblique shock of angle β_2 is drawn backwards

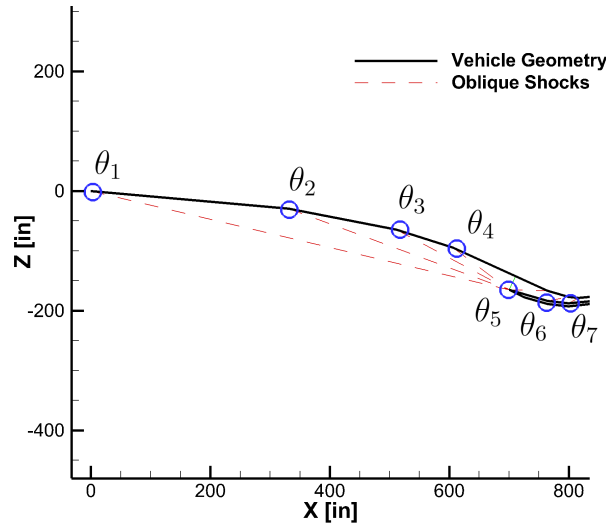


Fig. 2 Inlet with four external compression ramps and three internal compression ramps.

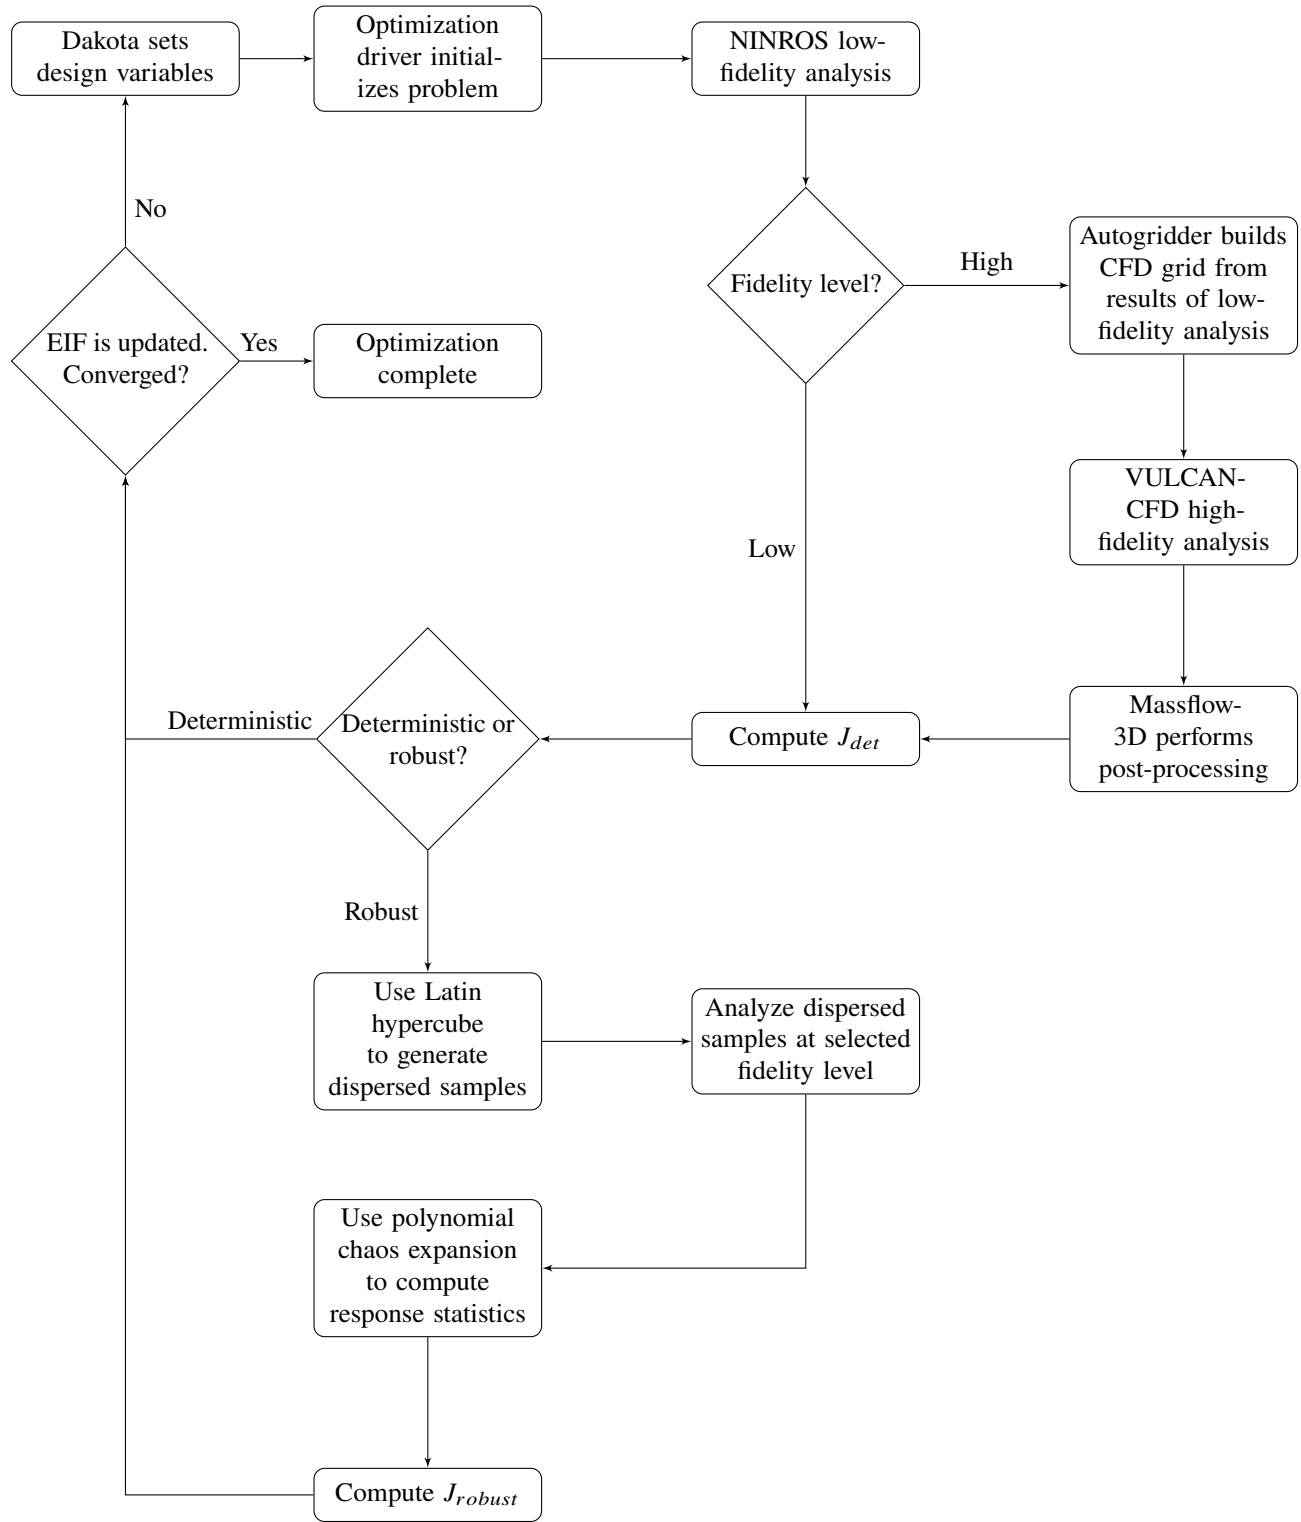


Fig. 3 Flowchart of the inlet design under uncertainty process.

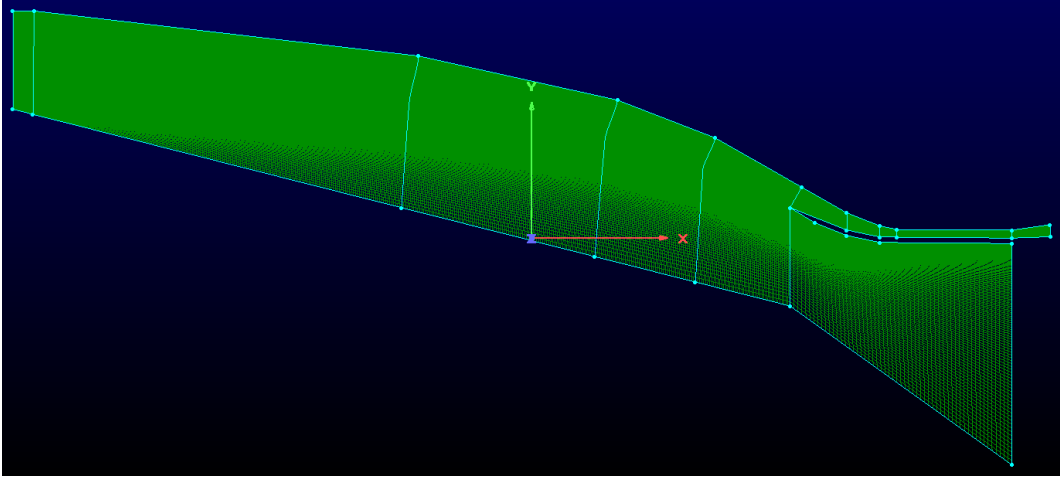


Fig. 4 Sample inlet grid topology.

from the cowl leading edge until intersecting with the ramp of θ_1 after the first shock. This intersection point is the location of the second ramp angle. The coordinates of the remaining external compression angles are found using this process of working backwards from the shock-on-lip condition. The result is a series of external compression ramps whose oblique shocks are all impinging upon the cowl leading edge, as shown previously in Figure 2.

The internal compression angles are determined in a slightly modified way. The angle of the cowl leading edge with respect to the inlet entrance plane flow, θ_{cle} , is known explicitly by the design variable corresponding to the first internal compression angle, in this case θ_5 . As in the external compression analysis, the oblique shock wave angle emitted from the cowl leading edge, β_{cle} , can be calculated, and its intersection point with the body side of the vehicle is computed. To ensure that the inlet does not have any reflected or intersecting oblique shocks, the wave of β_{cle} must be cancelled at the body-side surface. This is done by placing an expansion corner at the intersection point between β_{cle} and the body-side surface. If the expansion corner turns the flow by the same amount as θ_{cle} , then there will be no reflected oblique shock, nor an emitted expansion wave. The next internal compression corner is placed at the same x coordinate as the expansion corner cancelling the previous oblique shock, but on the cowl-side of the inlet. As with the external compression corners, this process is repeated until all compression and expansion corners have been placed. The Mach number at the throat and the total pressure recovery at the throat can be calculated by stepping through the oblique shock calculations, and the resulting geometry data is saved.

After the low-fidelity analysis executes, the analysis wrapper reaches a branch point. If low-fidelity mode is selected, then the deterministic objective function is calculated using Eq. (26a). If high-fidelity mode is enabled, then the analysis wrapper feeds the geometry data produced from the low-fidelity analysis to an autogridding routine. This autogridder routine controls the Pointwise software to build a viscous CFD mesh of the specified geometry. A sample grid that results from this process is shown in Figure 4. The mesh is exported, and then VULCAN-CFD is executed on the newly built grid. The Massflow-3D post-processing routine executes after VULCAN-CFD completes and produces one-dimensionalized mass-weighted averages of flow properties through the inlet as a function of x . The mass-weighted averages of total pressure recovery and Mach number at the throat are exported back to the analysis wrapper, which uses them to compute the deterministic objective function J_{Det} .

Once the deterministic objective function has been computed, regardless of fidelity level selected, another branch point is reached. If deterministic optimization is enabled, then the optimization driver feeds the deterministic objective function back into Dakota and convergence is checked. If robust optimization is enabled, then a Latin hypercube sampling (LHS) [32] method is used to generate dispersed samples of the uncertain parameters according to Table 2. The dispersed samples are fed back to the analysis wrapper, which runs the additional cases at the selected fidelity level. To maximize efficiency, all dispersed samples are run in parallel with one another. Once all UQ cases have finished, the optimization driver then uses polynomial chaos expansions to compute the mean and standard deviation of responses from the dispersed samples. With the response statistics known, the robust objective function is calculated using Eq. (28a), and J_{robust} is passed back into Dakota. The EIF is updated, and if the maximum EIF is $< 1\%$ of the current best value, the optimization is considered converged and terminates. If not, then another optimization iteration

is executed.

IV. Results

In this section, optimization results and analysis for several studies are presented. First, Section IV.A presents results from low-fidelity deterministic sweeps through the deterministic weighting coefficients. Both the EGO and the SQP algorithms are used to estimate numerical uncertainty in the global surrogate-based derivative-free approach. The numerical uncertainty estimated from this analysis is critically important and is used throughout the following studies. Next, Section IV.B presents high-fidelity results for the purpose of understanding how viscous effects influence inlet optimization. Section IV.C presents uncertainty propagation results using several different fidelity analyses. This analysis establishes the credibility of using multifidelity UQ techniques to closely approximate high-fidelity UQ. Section IV.D presents robust optimization results using both low-fidelity and multifidelity UQ approaches. Finally, Section IV.E discusses the physical cause of robustness in inlet designs and its link to fundamental physical principles.

A. Low-fidelity Deterministic Study

The first study conducted is a low-fidelity deterministic sweep through deterministic weighting coefficients such that $W_M + W_\Pi = 1$. The optimization problem being solved consists of Eqs. (26a) through (26c). The primary purpose of this study is to reveal the tradeoffs that can be made between the two inlet design objectives: M_{throat} and Π_{throat} . The secondary purpose of this study is to solve the optimization problem using both the global surrogate-based derivative-free EGO and gradient-based SQP algorithms. This allows for estimation of numerical error and uncertainty bands which are applied to EGO results throughout the following studies.

Results from the low-fidelity deterministic study are shown in Figure 5. Figure 5a shows the optimal design variables returned by SQP as a function of W_M . The SQP results show a smoothly varying trend that as W_M increases, each design variable grows larger. This is because as M_{throat} is weighted more heavily, the optimizer must drive this objective closer to its target value. Conversely, as M_{throat} is weighted less, the optimizer obtains slack in that objective which is used to increase the performance of Π_{throat} . Smaller total pressure losses are caused by weaker oblique shocks which result from smaller angles. Thus, small values of W_M will result in smaller ramp angles, while larger values of W_M force the ramp angles to increase. This trend is not smooth at the upper and lower bounds of the weighting coefficients, caused by an abrupt change in the objective function due to one of the two terms entirely dropping out. In addition to the trend seen as W_M increases, another trend is visible. For a given constant W_M , each successive ramp angle is slightly larger than the previous angle. This is the trend noted by Oswatitsch [1] which results in constant strength shocks. Constant strength shocks across a series of ramps results in minimizing the total pressure losses. This trend is visible at all weighting coefficients except $W_M = 1.0$ because this trend is the most efficient way for a 2D planar inlet to reach a given Mach number. The reason the Oswatitsch trend is not visible at $W_M = 1.0$ is because total pressure recovery entirely drops from the objective function at this weighting coefficient, thus the optimizer returns a configuration with a throat Mach number as close to the target value as possible with no regard for thermodynamic efficiency.

Figure 5b shows the optimal design variables that are returned by the EGO algorithm. Unlike the SQP results, the EGO results are noisy. Neither the smooth growth in design variables trend nor the Oswatitsch trend are clearly discernible. Comparing Figures 5a and 5b demonstrates the importance of quantifying the numerical error present in EGO results for this particular problem. Using Eq. (24), the optimizer early exit error in the design variables was found as a function of W_M . The minimum early exit error occurs at $W_M = 0.5$, which is used as the deterministic weighting coefficient throughout the rest of this work. The numerical uncertainty due to early exit error at $W_M = 0.5$ is calculated using Eq. (25) and is presented in Table 3 below.

Table 3 Estimated numerical uncertainty due to optimizer early exit.

	θ_i (°)	M_{throat}	Π_{throat}	Total turning (°)
U_{ee}	± 1.117	± 0.0106	± 0.0012	± 0.2098

Figures 5c and 5d show the objectives for the deterministic optimization problem, M_{throat} and Π_{throat} respectively. Results from SQP and EGO are plotted on the same figures to allow for easier comparison between the two. While there is still some noise present in the objectives returned from EGO, there is a much closer match to SQP results than is seen in the design variables. The major trends in the objectives are that as W_M increases, the value of M_{throat}

gets closer to its target value of $M_{target} = 3.0$, and as W_M decreases, then Π_{throat} gets closer to its target value of $\Pi_{target} = 1.0$. The optimizer early exit error is calculated and the numerical uncertainty for $W_M = 0.5$ is estimated. Results are given in Table 3. Figure 5e shows the Pareto frontier, a curve in objective-space indicating the optimal designs. The Pareto frontier clearly illustrates the tradeoffs available to the design engineer. Due to low error in the EGO results for objectives, there are no major differences between the Pareto frontiers plotted for the two algorithms.


Related to the design variables is the total turning; the sum of all design variables. This quantity is shown in Figure 5f for both EGO and SQP as a function of W_M . Similarly to the objectives, the EGO results for total turning are less smooth than the SQP results but still show considerable improvement over the design variables. The trend of increasing W_M leading to larger ramp angles is clearly visible in both the SQP and the EGO results. As with previous quantities, the optimizer early exit error is calculated and its resulting numerical uncertainty is estimated. Results are reported in Table 3. Note that as an integrated quantity, error in design variables tends to cancel out when summed up to the total turning. This canceling of error is what causes the EGO curve to be smoother for total turning than it is for the design variables. Accordingly, the numerical uncertainty estimate in total turning is roughly five times lower than it is for the design variables. This indicates that when trying to observe trends or interpret results from the Dakota implementation of EGO, the total turning is likely to be a more reliable source of information than the raw design variables.

The final result for the low-fidelity deterministic study is a comparison of the geometries from the extreme ends of the weighting coefficient scale. Figure 6 shows the inlet visualizations for a highly Π_{throat} -focused design and a highly M_{throat} -focused design. Observe that the Π_{throat} -focused design has smaller ramp angles to produce weaker shocks, as expected from discussion on design variables, while the M_{throat} -focused design has larger angles to produce the compression required to drive M_{throat} closer to its target value. A side effect of the weighting coefficients, not clearly discernible from any of the numerical results shown in Figure 5, is that lower values of W_M tend to elongate the inlet, even though all designs have a fixed cowl leading edge x coordinate. This is because shallower shock angles emitted from the cowl side of the inlet take more distance to reach the body side of the inlet. The shock-canceling design strategy requires placing expansion corners at the location of shock impingement on the body side, and successive compression corners are placed on the cowl side directly below the body side expansion corners. Thus, shallower angles stretch out the inlet. Accordingly, inlet designs that are more Π_{throat} -focused should be expected to weigh more and require larger vehicles for successful integration. Due to the elongation, and thus longer surface for skin friction to apply over, Π_{throat} -focused designs resulting from inviscid optimization may be expected to suffer from greater viscous losses than M_{throat} -focused designs when evaluated with high-fidelity analysis.


B. High-fidelity Deterministic Study

The second study conducted is a high-fidelity study. This study has two components. First, the optimal design from the low-fidelity deterministic optimization with $W_M = W_\Pi = 0.5$ from Section IV.A is evaluated at high-fidelity. The purpose of this single CFD evaluation is to demonstrate the differences in objective values that are produced for the same geometry when evaluated using different fidelity levels. The second component is running a deterministic optimization at high-fidelity with $W_M = W_\Pi = 0.5$. The purpose of this optimization is to understand how optimal inlet designs change when viscous effects are considered. Due to the substantially greater cost of this optimization than those conducted in Section IV.A, it is only feasible to complete a single run, rather than sweeping through multiple weighting coefficient values to obtain a Pareto frontier.


The objective values for the low-fidelity deterministic optimized geometry with $W_M = W_\Pi = 0.5$ are reported below in Table 4 for both low-fidelity analysis and high-fidelity analysis. For M_{throat} , the low-fidelity analysis gives roughly 8% error relative to the high-fidelity analysis. For Π_{throat} , the low-fidelity analysis gives roughly 42% error relative to the high-fidelity analysis. These discrepancies can be explained by examining the CFD flow contours in Figure 7. Mach number contours are shown in Figure 7a, while total pressure recovery contours are shown in Figure 7b. Figure 7a shows that the CFD results agree well with the low-fidelity prediction of M_{throat} , roughly 2.9 compared to 3.1. Figure 7b shows that the low-fidelity analysis is accurately predicting the total pressure recovery of the core flow within the inlet, roughly 91%, but the boundary layers on both the body side and cowl side of the inlet are imposing substantial losses. The two boundary layers consist of roughly 30% to 40% of the inlet duct at the throat, and much of the boundary layers have a total pressure recovery value from roughly 20% to 50%. These viscous losses, unpredicted by the inviscid low-fidelity analysis, are the source of the error shown in Table 4.




(a) Optimal design variables from SQP.




(b) Optimal design variables from EGO.




(c) M_{throat} objective values.



(d) Π_{throat} objective values.



(e) The Pareto frontier.



(f) Total turning.

Fig. 5 Optimization results from the low-fidelity deterministic study as a function of W_M , with $W_M + W_\Pi = 1$, using both SQP and EGO.

Table 4 Inlet throat conditions for the low-fidelity deterministic optimal design with $M_M = W_\Pi = 0.5$.

	M_{throat}	Π_{throat}
Low-fidelity model	3.090	0.9111
High-fidelity model	2.852	0.6429
Relative error in low-fidelity model	8.35%	41.7%

The next component of the high-fidelity study consists of deterministic optimization results for $W_M = W_\Pi = 0.5$. Results for the two objectives and total turning are compared to the low-fidelity optimal results in Table 5 below. Note that the high-fidelity optimal values have epistemic uncertainty as estimated from the early exit analysis of Section IV.A and reported in Table 3. Observe that due to the large viscous losses that can be imposed on Π_{throat} in the high-fidelity analysis, the optimizer is driven to a solution with substantially less total turning than for the same weighting coefficients at low-fidelity. This tends to penalize performance in M_{throat} , as the lessened total turning results in less compression. Figure 8 shows a comparison of the design variables that result from deterministic optimization at the two fidelity levels. The low-fidelity design variables are determined from the SQP algorithm, and thus do not have error bars on them. The SQP-returned values are taken to be the true optimum values for a given optimization setting. The high-fidelity design variables, returned from EGO, have epistemic uncertainty associated with them due to optimizer early exit. Although the first six design variables display the Oswatitsch trend, the epistemic uncertainty allows for the possibility that the Oswatitsch trend is not maintained. Interestingly, the final design variable shows a significant departure from the Oswatitsch trend. The magnitude of this departure is larger than the estimated numerical uncertainty due to EGO early exit, thus the departure should be assumed to be due to viscous effects. Possible reasons for a decrease in the final design variable, deepest within the inlet and where the boundary layers are at their thickest, are to minimize the likelihood of boundary layer separation and shock-boundary layer interactions. This observation is a topic of future research using optimization algorithms with less numerical uncertainty in design variables.

C. Uncertainty Propagation Study

The next study conducted is propagation of input uncertainties from Table 2 to output uncertainties in the objectives. This study is conducted on the geometry that results from the low-fidelity deterministic optimization of Section IV.A with $W_M = W_\Pi = 0.5$. The primary purpose of this study is to establish how closely multifidelity UQ can approximate purely high-fidelity UQ for the inlet design problem. This is because, as discussed in Sections II.C and II.D, robust optimization can be exceedingly computationally expensive. Thus, there is a strong incentive to use a multifidelity UQ approach rather than a purely high-fidelity approach when the two yield comparable results. The secondary purpose of this study is to understand the differences in uncertainty captured by different analysis fidelity levels.

Figure 9a shows the results of propagating input uncertainties on M_{throat} , while 9b shows results for Π_{throat} . For both M_{throat} and Π_{throat} the multifidelity UQ approach shows excellent agreement with the purely high-fidelity UQ approach for this geometry. This agreement suggests that when robust optimization is performed, multifidelity UQ




Fig. 6 Comparison of the optimal geometries resulting from low-fidelity deterministic optimization.

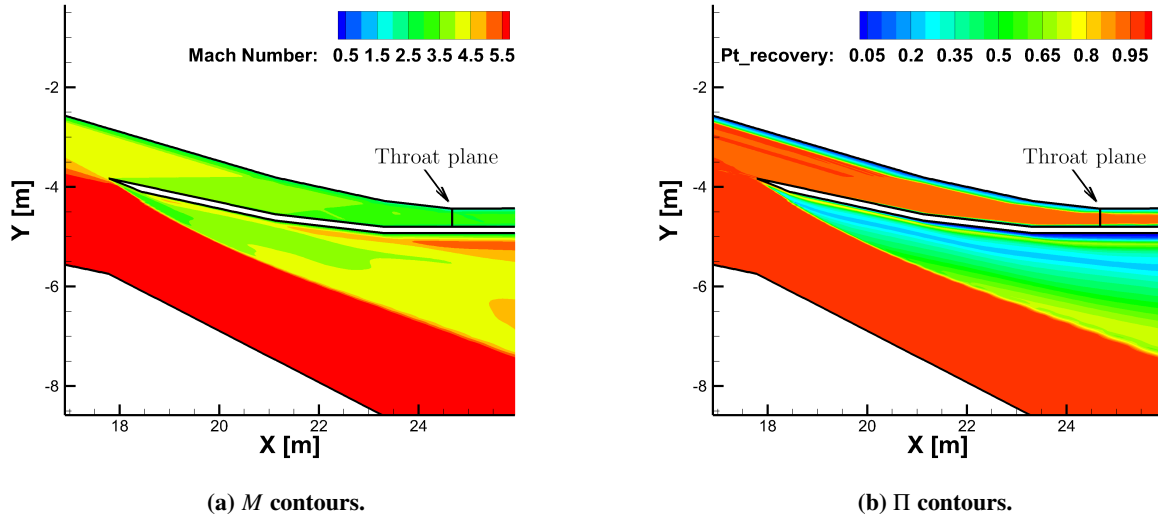


Fig. 7 High-fidelity CFD contours of the $W_M = W_\Pi = 0.5$ low-fidelity deterministic optimal design.




Fig. 8 Design variable comparison between low-fidelity and high-fidelity deterministic optimum designs.

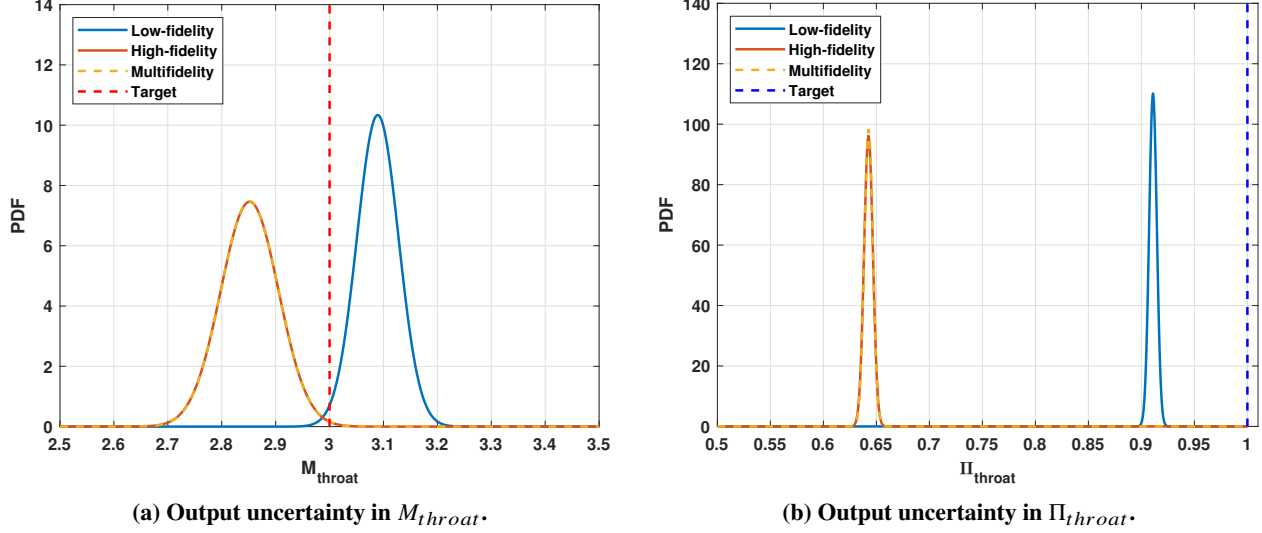


Fig. 9 Uncertainty propagation results for the low-fidelity deterministic optimal design with $W_M = W_\Pi = 0.5$.

can be used in place of purely high-fidelity UQ, assuming that agreement is similar among all geometries analyzed. Referring to Eq. (18) with an oversampling ratio of 1.5 and $n = 3$ random dimensions, reducing polynomial order q from 2 to 1 yields 6 high-fidelity cases for the multifidelity approach, compared to 15 for the purely high-fidelity approach. This reduces the cost of robust optimization to 40% of what it would be if purely high-fidelity UQ was performed.

The mean and standard deviation for both objectives at all three fidelity levels are reported in Table 6 below. Note that low-fidelity UQ does not capture all of the uncertainty. For M_{throat} , the low-fidelity uncertainty is roughly 72% of that for high and multifidelity analysis. For Π_{throat} , the low-fidelity analysis captures roughly 87% of the uncertainty compared to the high and multifidelity analyses.

Table 6 Results from the uncertainty propagation study.

	μ_M	σ_M	μ_Π	σ_Π
Low-fidelity	3.089	0.03860	0.9109	0.003622
High-fidelity	2.852	0.05347	0.6426	0.004148
Multifidelity	2.852	0.05336	0.6427	0.004058
Error in low-fidelity relative to high-fidelity	8.31%	27.8%	41.8%	12.7%
Error in multifidelity relative to high-fidelity	0.00%	0.206%	0.0156%	2.17%

D. Robust Optimization Study

This section presents results from robust optimization using both low-fidelity and multifidelity approaches. The optimization problem solved is now that of Eqs. (28a) through (28c). The robust optimization study ties in results from the early exit uncertainty analysis of Section IV.A, viscous effects from Section IV.B, and UQ results from Section IV.C. First, a low-fidelity sweep through the robust weighting coefficients is performed such that $W_\mu + W_\sigma = 1$. The deterministic weighting coefficients are held constant at $W_M = W_\Pi = 0.5$. The purpose of the robust weighting

Table 5 Results from the low-fidelity optimization compared to high-fidelity optimization for $W_M = W_\Pi = 0.5$.

	M_{throat}	Π_{throat}	Total turning ($^\circ$)
Low-fidelity optimum	3.090	0.9111	32.69
High-fidelity optimum	3.788 ± 0.0106	0.9501 ± 0.0012	21.69 ± 0.2098

coefficient sweep is to understand the design tradeoffs that can be made as either mean performance or variance of performance is prioritized more. Once the robust weighting coefficient sweep is complete, a single set of coefficients is selected for more detailed comparison to low-fidelity deterministic optimization. The purpose of this comparison is to more thoroughly demonstrate the design changes incurred by robust design. Finally, the results of multifidelity robust optimization are presented for the same weighting coefficients as used in the previous comparison. The purpose of this comparison is to demonstrate the effects of viscous boundary layers on robust optimization.


Results from the sweep through robust weighting coefficients are shown in Figure 10. The optimal design variables returned from EGO are shown in Figure 10a and the total turning and net turning are shown in Figure 10b. As seen previously in Section IV.A, the raw design variables have significant noise, but the integrated quantity of total turning is smooth. The total turning is observed to be roughly flat at intermediate and large values of W_μ , but starts to substantially increase as W_μ decreases. This indicates that more total turning is associated with decreased variance in performance. Figures 10c and 10d show the mean performance of the objectives, μ_M and μ_Π , respectively. As W_μ becomes small, and therefore as W_σ becomes large, the mean performance of both objectives falls precipitously. Figures 10e and 10f show the standard deviation of the two objectives, σ_M and σ_Π . As the optimizer increases W_σ , the standard deviation of M_{throat} is seen to drop to roughly 60% of its value when W_μ is near 1. However, the standard deviation of Π_{throat} more than doubles from its value when W_μ is near 1. The increase in σ_Π , opposite the desired trend, is caused by the differing magnitudes of σ_M and σ_Π . While σ_Π is increasing as W_σ increases, the simultaneous larger decrease in σ_M more than makes up for it. These results indicate that there are opposing trends in the inlet design qualities that minimize σ_M and σ_Π . Note that due to the form of the robust objective function in Eq. (28a), the deterministic weighting coefficients still influence which throat objective is prioritized more in robust optimization. Thus, adjusting the deterministic weighting coefficients gives the design engineer the ability to alter the behavior of the robust objective function to minimize σ_Π rather than σ_M .

Figure 11 shows four different Pareto frontiers that can be formulated from the robust objective function of Eq. (28a). Figure 11a shows the Pareto frontier between both μ values. The μ frontier indicates that there is not a noteworthy tradeoff between μ_M and μ_Π , as the design qualities that improve one also improve the other. Figure 11b shows the Pareto frontier between both σ values. The σ frontier indicates that, unlike μ , a design tradeoff must be made between σ_M and σ_Π because they are inversely related to one another. Figure 11c shows the Pareto frontier for M_{throat} performance, which indicates that a tradeoff must be made between μ_M and σ_M . Finally, Figure 11d shows the Pareto frontier for Π_{throat} performance. The Π_{throat} frontier indicates that there is no tradeoff to be made between μ_Π and σ_Π because the design qualities that improve one quantity also improve the other.

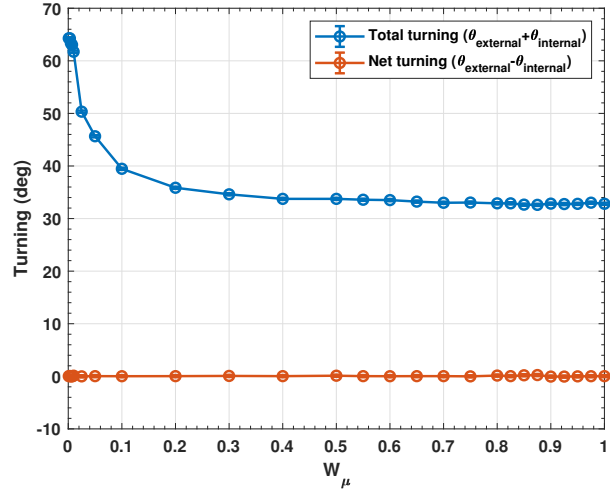
Synthesizing observations from the plots of total turning, objectives, and Pareto frontiers, conclusions are drawn regarding the design trends incurred by robust optimization. The Pareto frontiers show that improvements in three of the four robust objectives, μ_M , μ_Π , and σ_Π , are related to one another. The final robust objective, σ_M , is inversely related to the other three. Due to differing magnitudes of σ_M and σ_Π , the set of deterministic weighting coefficients chosen for this study results in robust optimization favoring improvements in σ_M over improvements in σ_Π . Improvements in σ_M are caused by increasing the total turning beyond the amount returned by deterministic optimization with $W_M = 1.0$. Thus, the variance in M_{throat} is minimized by over-compressing the flow to beyond the throat target value M_{target} , driving μ_M further from its target. The excessive compression caused by the increase in total turning requires stronger shocks, leading to greater losses, therefore decreasing μ_Π . The above discussion gives insight into how the four robust objectives are linked to one another and how robust optimization is affecting the mean values. However, it does not reveal why the total turning is linked to changes in the variance of throat objectives. This topic is investigated in more detail in Section IV.E.

The next component of the robust optimization study is a comparison with low-fidelity deterministic results from SQP, shown in Figure 12. The total turning for low-fidelity robust optimization compared to low-fidelity deterministic optimization is shown in Figure 12a. The deterministic total turning is plotted against W_M , while the robust total turning is plotted against W_μ . All of the robust optimization results have $W_M = W_\Pi = 0.5$. This plot shows that, as discussed previously, decreasing W_μ and increasing W_σ causes the total turning to roughly double its value compared to when W_μ is near 1.0. Figure 12b shows the same quantities as Figure 12a but zoomed in on the $0.5 \leq W_\mu \leq 1.0$ region. Figure 12b illustrates that as the robust weighting coefficient W_μ approaches 1.0, the value of the total turning returned by robust optimization should approach the value returned by deterministic optimization with the same deterministic weighting coefficients.


Figures 12c and 12d show the probability density functions (PDFs) of the throat objectives for the deterministic optimum compared to the robust optimum with $W_\mu = 0.1$ and $W_\sigma = 0.9$. Robust optimization for high W_σ values is overly compressing the flow, reducing the throat Mach number beyond its target value. Excess compression results



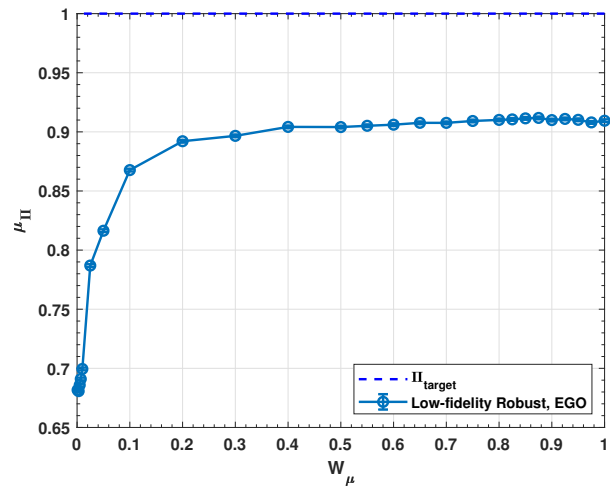
(a) Design variables.




(b) Total turning and net turning.



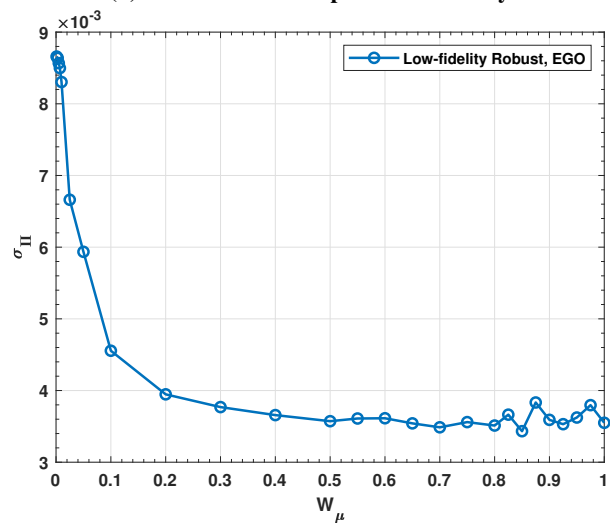
(c) Mean throat Mach number.



(d) Mean throat total pressure recovery.



(e) Standard deviation of throat Mach number.



(f) Standard deviation of throat total pressure recovery.

Fig. 10 Optimization results from the low-fidelity robust study with $W_M = W_{\Pi} = 0.5$ and $W_{\mu} + W_{\sigma} = 1.0$.

in stronger shocks, which tends to decrease μ_{Π} . The standard deviation σ_M decreases by roughly 15% from robust optimization, while σ_{Π} increases by roughly 26%. A comparison of the objective statistics from these two optimizations is presented in Table 7 below.

Table 7 Comparison of results for low-fidelity deterministic to robust.

	μ_M	σ_M	μ_{Π}	σ_{Π}
Low-fidelity Deterministic	3.089	0.03860	0.9109	0.003622
Low-fidelity Robust	2.726	0.03294	0.8677	0.004555

Figure 12e displays the design variables from robust optimization compared to deterministic optimization. While the angles are generally larger across the inlet for robust optimization, in accordance with the observed increase in total turning, there appear to be no departures from the Oswatitsch trend that can not be explained by epistemic uncertainty.

The final component of the robust optimization study is a multifidelity robust optimization. Similar to the high-fidelity deterministic optimization of Section IV.B, running more than one multifidelity robust optimization is infeasible due to

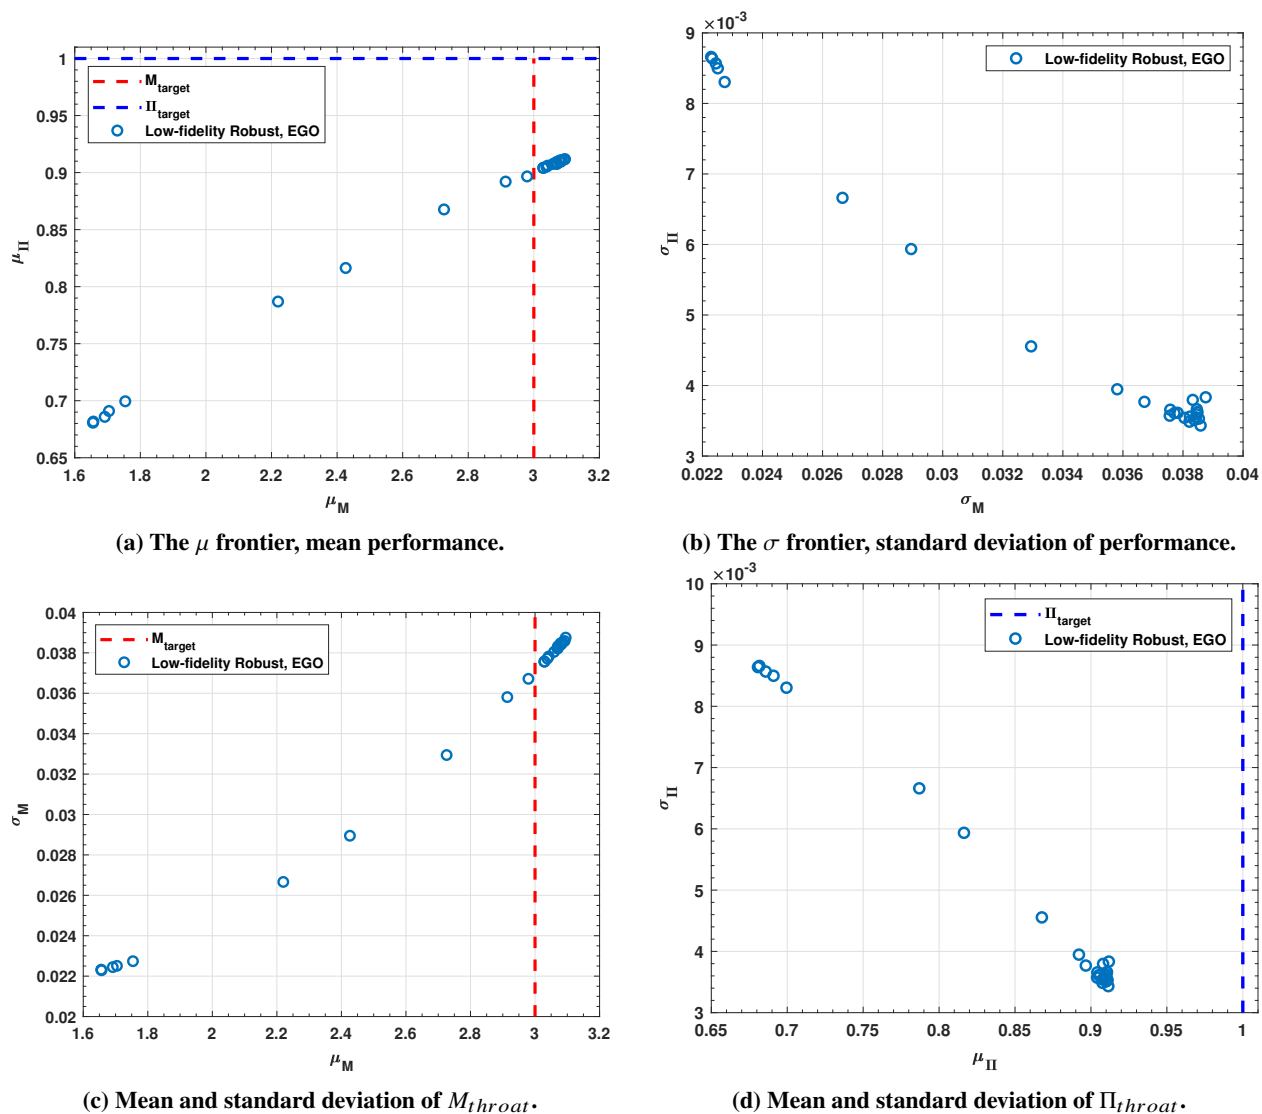

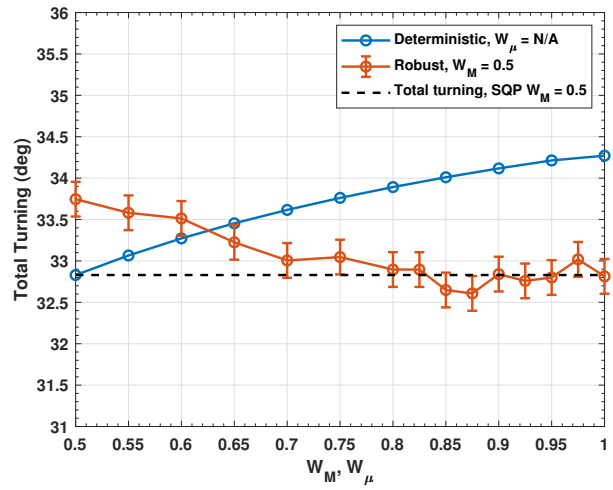


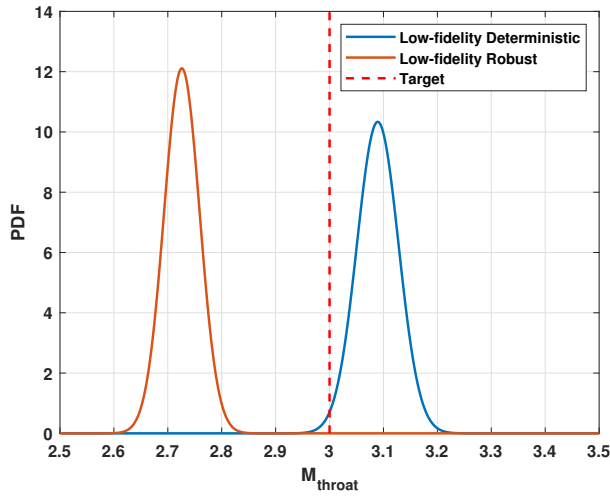
Fig. 11 Pareto frontiers from the low-fidelity robust study with $W_M = W_{\Pi} = 0.5$ and $W_{\mu} + W_{\sigma} = 1.0$.



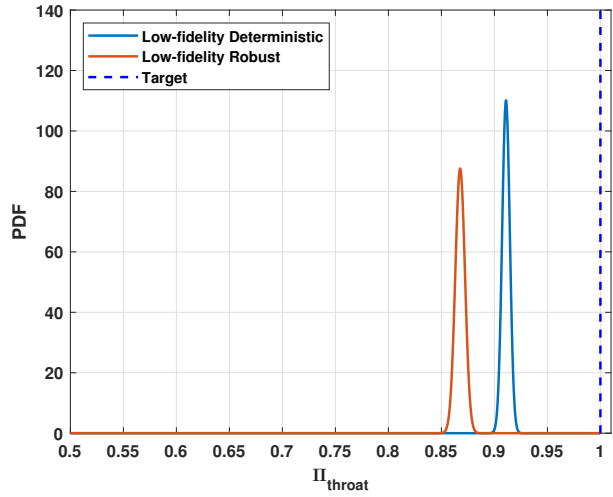
(a) Total turning.



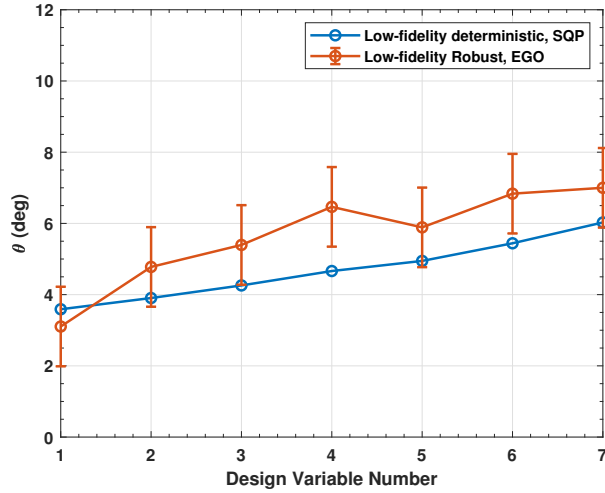
(b) Total turning zoomed in on the $W_M = W_{\Pi} = 0.5$.



(c) PDFs for M_{throat}



(d) PDFs for Π_{throat}



(e) Design variables.

Fig. 12 Low-fidelity robust compared to low-fidelity deterministic results.

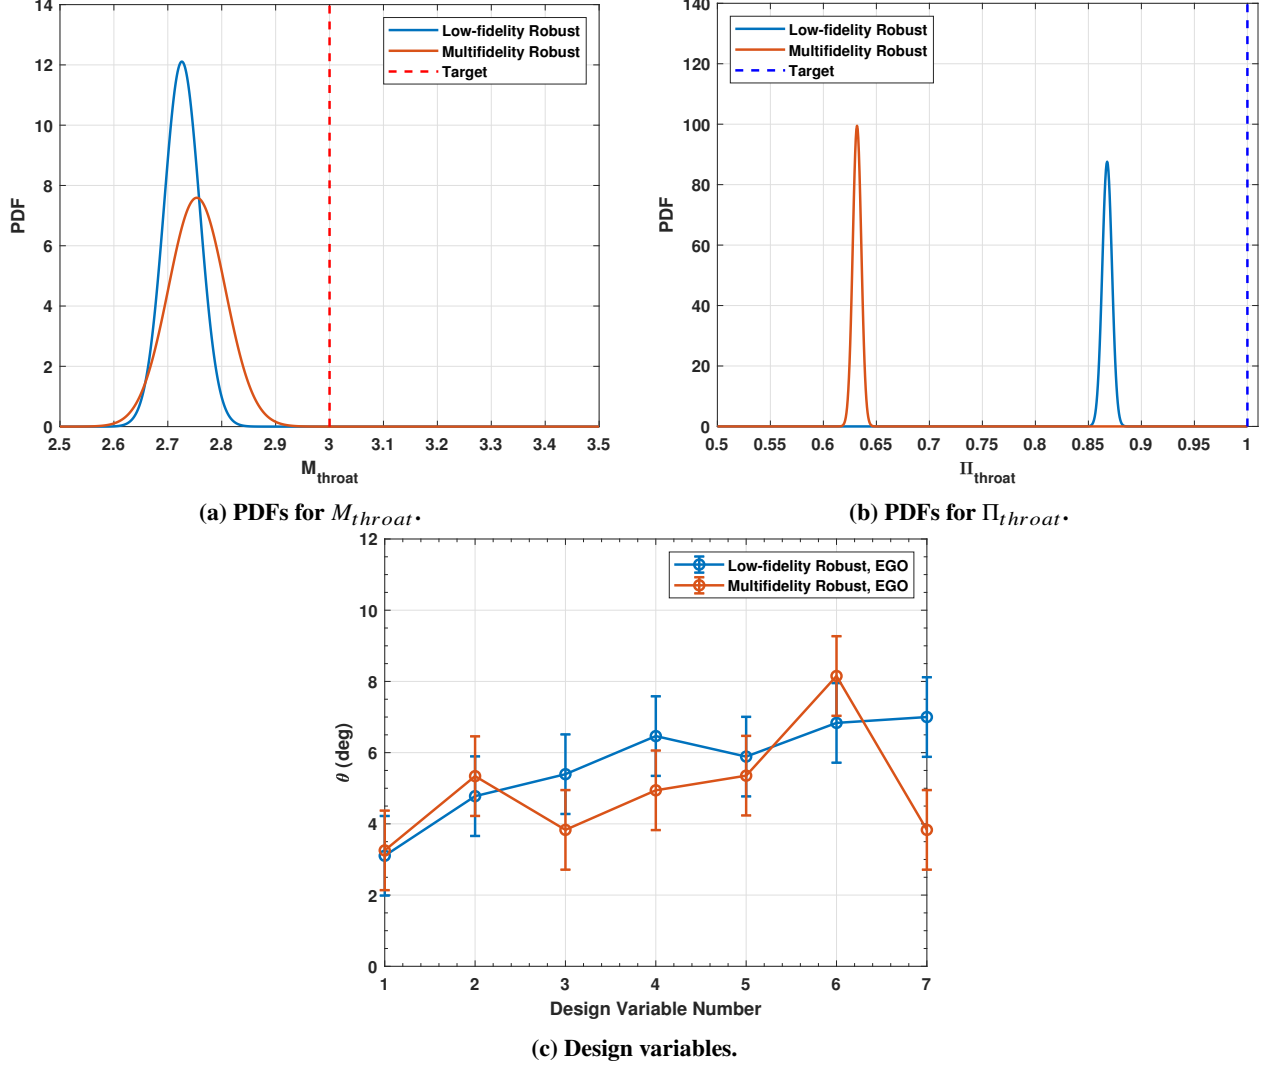


Fig. 13 Multifidelity robust compared to low-fidelity robust results.

computational expense. As a result there are no obtained Pareto frontiers to give insight on the tradeoffs available. The optimization is run with $W_M = W_\Pi = 0.5$, $W_\mu = 0.1$, and $W_\sigma = 0.9$. The PDFs of the throat objectives are shown in Figures 13a and 13b. The objective distributions indicate that compared to low-fidelity robust optimization, multifidelity robust optimization is predicting roughly 60% more uncertainty in M_{throat} and roughly 12% less uncertainty in Π_{throat} . The design variables are shown in Figure 13c. As with the high-fidelity deterministic optimization of Section IV.B, there is once again a departure from the Oswatitsch trend in the final design variable that is larger than the estimated numerical uncertainty from EGO early exit. There appears to be a departure from the Oswatitsch trend in design variables two and three as well, but this is within epistemic uncertainty bounds so no strong conclusions can be drawn about this observation. As with the results of high-fidelity optimization, this departure from the Oswatitsch trend when viscous effects are considered is a topic for future research.

E. Physical Cause of Robustness

The goal of high quality optimization work should not be to simply run a computer program and accept the results, but to use an optimizer as a tool to aid in understanding physical trends. In this vein, the final study of the present work is an investigation into the underlying physical cause that leads to some inlet configurations being robust, while others are sensitive.


The θ - β - M function of Eq. (1), or equivalently the β - θ - M function of Eqs. (2) through (4), provides the oblique shock wave angle β for a known pre-shock Mach number M_1 and deflection angle θ . The values of M_1 , θ , and β specify the post-shock total pressure recovery Π and Mach number M_2 , shown in Figures 14a and 14b for several sample Mach numbers across a range of deflection angles. These curves all have varying slope and curvature across their length. Understanding these variations leads to an explanation of why some designs are robust and others are sensitive. The sensitivity of the output quantity, the y-axis value, on these plots is related to the derivative with respect to the x-axis value. As an example, consider the Mach 10 curve on the Π plot of Figure 14a. If there is a deflection angle with some small uncertainty in the precise value, $\theta \pm \delta$, centered about $\theta = 10^\circ$, then there will be substantial uncertainty in the response quantity Π due to how steep the Mach 10 curve is at $\theta = 10^\circ$. Therefore, this configuration is sensitive in Π . In contrast, if the uncertain deflection angle $\theta \pm \delta$ is centered about $\theta = 1^\circ$ with the same amount of uncertainty, then there is very little uncertainty in the output Π due to how flat the slope is at $\theta = 1^\circ$. Therefore, this configuration is robust in Π . Both of these examples use the same value of uncertainty in the input condition, δ , but the sensitivity of the system changes based on the derivative at the current conditions. Examining the rest of the Mach curves on Figure 14a reveals that they all tend to have small slopes as the deflection angle approaches 0, and steeper slopes in roughly the $10^\circ \leq \theta \leq 30^\circ$ region. Particularly large Mach numbers, such as Mach 10, have an additional low-slope robust region around roughly $\theta \geq 30^\circ$, but the low total pressure recovery obtained from such a configuration renders this irrelevant for airbreathing propulsion purposes. Note that the observation that decreasing deflection, ideally as close as possible to 0, increases robustness in Π matches the result obtained in Section IV.D for total turning, shown in Figures 10b and 10f.

When the above analysis is repeated for the post-shock Mach number M_2 , the results are not as immediately clear. The curves for large Mach numbers, such as Mach 10, in Figure 14b are steeper at low deflection angles and flatter at intermediate deflection angles, but it is more difficult to interpret the slope for lower Mach number curves. To assist with this analysis, finite differences are used to calculate the derivatives of the post-shock quantities. The derivative of Π with respect to θ is shown in Figure 14c, and the derivative of M_2 is shown in Figure 14d, for deflection angles relevant to the bounds of the optimization problems in this work. On the derivative plots, robustness is indicated by derivative values close to zero. For Π , the trends visible are that deflection angles close to zero are more robust than higher deflections, and lower Mach numbers are more robust than higher Mach numbers. For M_2 , lower Mach numbers are once again seen to be more robust than higher Mach numbers. The trend for deflection angle is less clear on Figure 14d, so a change of variables is performed. Solutions are found for several sample deflection angles across a range of pre-shock Mach numbers, plotted in Figure 14e. The derivatives of these solutions are shown in Figure 14f. Here, the derivative plot shows that above some crossover Mach number around roughly 1.8, higher θ values decrease the absolute value of the derivative for M_2 . Therefore, to decrease sensitivity in M , the amount of deflection should increase. These results match the numerical trends observed in Section IV.D, where it was found that increases in total turning increased the robustness of M_{throat} by minimizing σ_M .

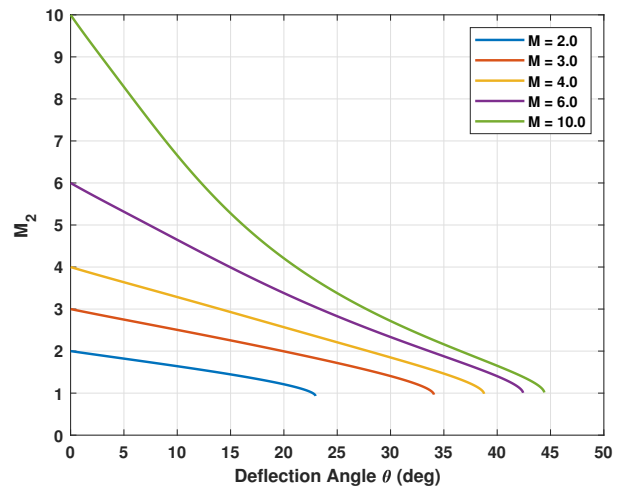
V. Conclusions

Scramjet inlet design studies have been done many times, typically either using inviscid analytical techniques, or deterministic methods. In this work, a scramjet inlet design optimization problem was formulated both deterministically and stochastically for both low and high-fidelity analyses. High-fidelity analysis consists of fully viscous CFD with grid-generation within the optimization loop. Robust optimization was performed using a global derivative-free optimization algorithm. A method of quantifying numerical uncertainty from the optimization algorithm was developed. The scramjet design under uncertainty problem was solved and trends in the design space were explored. Findings include that viscous effects may result in departure from the Oswatitsch trend, multifidelity uncertainty quantification nearly perfectly replicates the results of purely high-fidelity uncertainty quantification for the inlet design problem, and design trends that increase inlet robustness were identified. Finally, the cause of robustness in scramjet inlet designs was linked back to fundamental physical principles through the θ - β - M function. Thorough analysis of the θ - β - M function and oblique shock relations was demonstrated to enable accurate predictions of which design trends will increase or decrease system sensitivity.

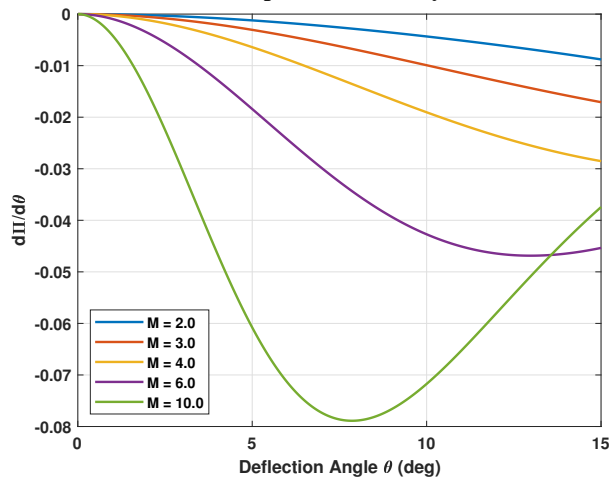
Future work on this topic will extend the current techniques to the problem of multipoint optimization under uncertainty for variable geometry inlets operating across a range of Mach numbers. Additional areas of investigation include using reliability-based design optimization to minimize the probability of failure events such as engine unstart, and optimizing combustor and nozzle designs under uncertainty. The long-term goal is to perform design under uncertainty at the overall system level for an airbreathing hypersonic vehicle.




(a) Total pressure recovery.




(b) Post-shock M .




(c) Total pressure recovery sensitivity.



(d) Post-shock M sensitivity.



(e) Post-shock M vs. pre-shock M .



(f) Sensitivity of post-shock M vs. pre-shock M .

Fig. 14 Manipulations of the $\theta - \beta - M$ function weak shock solutions.

References

- [1] Oswatitsch, K., “Pressure Recovery for Missiles with Reaction Propulsion at High Supersonic Speeds (The Efficiency of Shock Diffusers),” 1947.
- [2] Smart, M., “Optimization of Two-Dimensional Scramjet Inlets,” *Journal of Aircraft*, Vol. 36, No. 2, 1999, pp. 430–433.
- [3] Hasegawa, S., and Knight, D., “Numerical Analysis and Optimization of Two-Dimensional Hypersonic Inlets,” *42nd AIAA Aerospace Sciences Meeting and Exhibit*, 2004, p. 856.
- [4] Brown, M., Mudford, N., Neely, A., and Ray, T., “Robust Design Optimization of Two-Dimensional Scramjet Inlets,” *14th AIAA/AHI Space Planes and Hypersonic Systems and Technologies Conference*, 2006, p. 8140.
- [5] Anderson, J. D., *Modern Compressible Flow: with Historical Perspective*, McGraw-Hill, 2003.
- [6] Baurle, R. A., White, J. A., Drozda, T. G., and Norris, A. T., “VULCAN-CFD Theory Manual: Ver. 7.1.0,” 2020.
- [7] Wilcox, D. C., et al., *Turbulence Modeling for CFD*, Vol. 3, DCW Industries, Inc. La Canada, CA, 2006.
- [8] Edwards, J. R., “A Low-Diffusion Flux-Splitting Scheme for Navier-Stokes Calculations,” *Computers & Fluids*, Vol. 26, No. 6, 1997, pp. 635–659.
- [9] Van Leer, B., “Towards the Ultimate Conservative Difference Scheme. II. Monotonicity and Conservation Combined in a Second-Order Scheme,” *Journal of Computational Physics*, Vol. 14, No. 4, 1974, pp. 361–370.
- [10] Menter, F. R., et al., “Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications,” *AIAA journal*, Vol. 32, No. 8, 1994, pp. 1598–1605.
- [11] Chan, T. F., and Van der Vorst, H. A., “Approximate and Incomplete Factorizations,” *Parallel Numerical Algorithms*, Springer, 1997, pp. 167–202.
- [12] Inc., C. D. S., “Pointwise Software Manual,” <https://www.pointwise.com/doc/user-manual/>, 2021.
- [13] Meyer, B., “User Manual and Methodology for One-Dimensionalization Code: Massflow3d version 8.4,” Tech. rep., Technical Note 06-474, NASA Langley Contract NAS1-00135B, 2006.
- [14] Eldred, M. S., and Elman, H. C., “Design Under Uncertainty Employing Stochastic Expansion Methods,” *International Journal for Uncertainty Quantification*, Vol. 1, No. 2, 2011.
- [15] Oberkampf, W. L., and Roy, C. J., *Verification and Validation in Scientific Computing*, Cambridge University Press, 2010.
- [16] Cameron, R., and Martin, W., “The Orthogonal Development of Nonlinear Functionals in Series of Fourier-Hermite Functionals,” *Annals of Mathematics*, Vol. 48, 1947, pp. 385–392.
- [17] Smith, R. C., *Uncertainty Quantification: Theory, Implementation, and Applications*, SIAM, 2013.
- [18] Xiu, D., and Karniadakis, G. E., “The Wiener–Askey Polynomial Chaos for Stochastic Differential Equations,” *SIAM Journal on Scientific Computing*, Vol. 24, No. 2, 2002, pp. 619–644.
- [19] Ng, L. W.-T., and Eldred, M., “Multifidelity Uncertainty Quantification using Non-Intrusive Polynomial Chaos and Stochastic Collocation,” *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference 20th AIAA/ASME/AHS Adaptive Structures Conference 14th AIAA*, 2012, p. 1852.
- [20] West IV, T. K., and Phillips, B. D., “Multifidelity Uncertainty Quantification of a Commercial Supersonic Transport,” *Journal of Aircraft*, Vol. 57, No. 3, 2020, pp. 491–500.
- [21] Eldred, M., “Recent Advances in Non-Intrusive Polynomial Chaos and Stochastic Collocation Methods for Uncertainty Analysis and Design,” *50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference 17th AIAA/ASME/AHS Adaptive Structures Conference 11th AIAA No.*, 2009, p. 2274.
- [22] Ng, L. W., and Willcox, K. E., “Multifidelity Approaches for Optimization Under Uncertainty,” *International Journal for Numerical Methods in Engineering*, Vol. 100, No. 10, 2014, pp. 746–772.
- [23] Peherstorfer, B., Willcox, K., and Gunzburger, M., “Survey of Multifidelity Methods in Uncertainty Propagation, Inference, and Optimization,” *SIAM Review*, Vol. 60, No. 3, 2018, pp. 550–591.

- [24] Jones, D. R., Schonlau, M., and Welch, W. J., "Efficient Global Optimization of Expensive Black-box Functions," *Journal of Global Optimization*, Vol. 13, No. 4, 1998, pp. 455–492.
- [25] Huang, D., Allen, T. T., Notz, W. I., and Zheng, N., "Global Optimization of Stochastic Black-box Systems via Sequential Kriging Meta-models," *Journal of Global Optimization*, Vol. 34, No. 3, 2006, pp. 441–466.
- [26] Adams, B., Bohnhoff, W., Dalbey, K., Ebeida, M., Eddy, J., Eldred, M., Hooper, R., Hough, P., Hu, K., Jakeman, J., et al., "Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.13 User's Manual." Tech. rep., Sandia National Lab.(SNL-NM), Albuquerque, NM (United States), 2020.
- [27] Gablonsky, J., "DIRECT Version 2.0 User Guide," Tech. rep., North Carolina State University. Center for Research in Scientific Computation, 2001.
- [28] Arora, J., *Introduction to Optimum Design*, Elsevier, 2017.
- [29] Roache, P. J., "Quantification of Uncertainty in Computational Fluid Dynamics," *Annual Review of Fluid Mechanics*, Vol. 29, No. 1, 1997, pp. 123–160.
- [30] Vanderplaats, G. N., "Multidiscipline Design Optimization," 1988.
- [31] Ferlemann, P. G., "Variable-Geometry High-Speed Inlet Design and Analysis for the Aether Turbine-Based Combined-Cycle Vehicle," *38th JANNAF Airbreathing Propulsion Meeting*, Newport News, VA, 2020.
- [32] McKay, M. D., Beckman, R. J., and Conover, W. J., "A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code," *Technometrics*, Vol. 21, No. 2, 1979, pp. 239–245.