

This is a preprint of the following article,
Guodong Chen, Krzysztof J. Fidkowski, Discretization error control for constrained aerodynamic shape optimization, Journal of Computational Physics, 2019 (In press). DOI: [10.1016/j.jcp.2019.02.038](https://doi.org/10.1016/j.jcp.2019.02.038)
The published article may differ from this preprint.

Discretization error control for constrained aerodynamic shape optimization

Guodong Chen^{a,*}, Krzysztof J. Fidkowski^a

^a*Department of Aerospace Engineering, University of Michigan, MI 48109, United States*

Abstract

In this paper, we present a method to control the discretization error in constrained aerodynamic shape optimization problems using meshes adapted via adjoint-based error estimates. The optimization constraints may involve outputs that are not directly targeted for optimization, and hence also not for error estimation and mesh adaptation. However, discretization errors in these outputs often indirectly affect the calculation of the objective function. The proposed method takes this effect into account, so that the mesh is adapted and itself optimized to predict both objective outputs and constraint outputs with appropriate accuracy. We use unstructured mesh optimization to maximize accuracy of the results for a given number of degrees of freedom. The error estimates drive a multifidelity optimization process, preventing over-optimization on a coarse mesh and over-refinement on an undesired design. We demonstrate the accuracy and efficiency of our proposed method on several airfoil shape optimization problems governed by the compressible Navier-Stokes equations. The framework is expected to be even more important for optimization problems with complex systems, dramatic design changes, or high-accuracy requirements.

Keywords: Constrained aerodynamic optimization, Multifidelity optimization, Discretization error control, Adjoint-based error estimation, Output-based mesh adaptation, Unstructured mesh optimization

1. Introduction

Over the past decades, Computational Fluid Dynamics (CFD) has benefited greatly from increasing power of modern computers and highly-developed numerical methods, such that it has presently become an integral part of vehicle design in aerospace engineering. The capacity of design evaluation at almost arbitrary test conditions, as well as lower costs compared to wind tunnel tests, have made CFD an indispensable tool in the aerodynamic design process. Although high-fidelity CFD simulations are now commonly used in analysis and design, high-dimensional optimization problems still remain computationally taxing due to the large number of function evaluations required to reach the optimum. Much research has been done to reduce this cost, particularly in gradient-based methods, which, compared to their gradient-free counterparts, generally converge to the optimum with fewer function evaluations and thus are more efficient for aerodynamic design. In gradient-based methods, gradients with respect to each design parameter are needed at every optimization iteration, which requires an accurate and efficient method for estimating sensitivities. Methods such as finite differencing, complex-step derivative approximation [1], and algorithmic differentiation [2] depend on the

*Corresponding author

Email addresses: cgderic@umich.edu (Guodong Chen), kfid@umich.edu (Krzysztof J. Fidkowski)

number of design parameters, while the adjoint method [3, 4, 5, 6, 7] is largely independent of the number of design parameters.

A diagram of practical aerodynamic optimization with gradient-based methods is shown in Figure 1(a). Many error sources exist in CFD analysis, and these errors can pollute the results of the optimization. The errors include modeling errors when the real-world system is modelled with simplified assumptions, e.g. incompressible, inviscid flow, numerical errors during the discretization of the governing PDEs on a finite-dimensional space, and convergence errors when solving the discretized system of equations. Modeling errors inherent to the governing equations can be reduced by model validation or by choosing more complex models, but the estimation of these errors is not addressed in our present work. Instead, we aim to efficiently and robustly solve optimization problems in which the chosen model is assumed to be exact. In addition, we do not address iterative solution convergence errors, which we control to be many orders of magnitude smaller than the optimization tolerance through appropriate solver settings. Therefore, the numerical errors considered in this paper refer to the errors caused by the discretization of the continuous governing equations on a finite-dimensional space.

Discretization errors strongly affect the reliability of the optimization objective and sensitivity calculations, and hence the quality of the optimization results. These errors appear in every single design analysis during the optimization, and the effect becomes even more severe when the optimizer sequentially uses inaccurate data from the numerical simulations. Figure 1(b) shows a simple optimization problem governed by a one-dimensional advection-diffusion PDE. The numerical solution is obtained by solving the system on a coarse uniform spatial discretization. There is only one local minimum for the original continuous problem, while for the discrete numerical solution, we find a spurious local minimum caused purely by numerical errors induced by the discretization. Depending on the starting point, the optimizer may (a) get stuck in a spurious optimum created by discretization errors, or (b) work on the discretization errors rather than on the physics and converge to an incorrect optimum. In order to avoid these undesired behaviors, the discretization errors must be carefully controlled in the optimization.

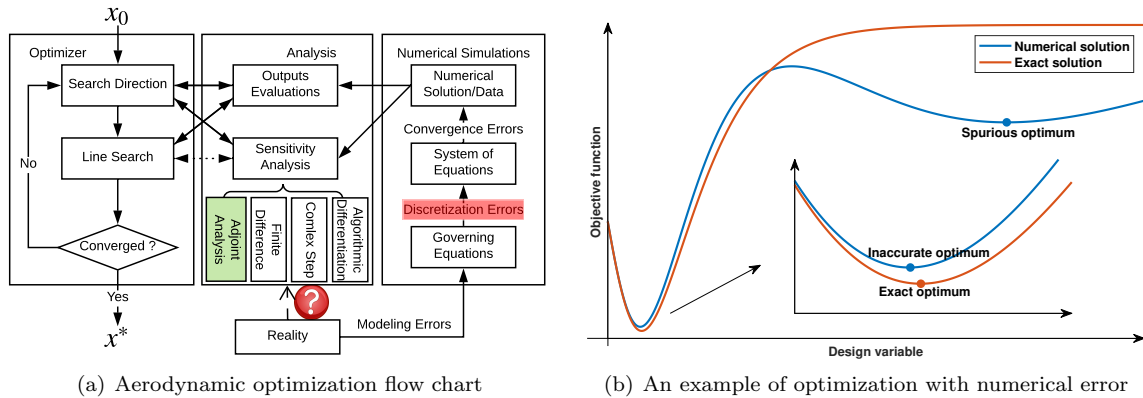


Figure 1: Effects of numerical error in optimization.

For a single fixed design, simulation accuracy and efficiency can be dramatically improved by adaptive methods, in which the discretization is iteratively improved through local mesh refinement and/or approximation order increment. Residual-based and feature-based error estimation and mesh adaptation can be robust for elliptic PDEs like those of structural elasticity [8, 9, 10, 11], and several contributions have been made to integrate these approaches to structure optimization problems [12, 13, 14]. However for hyperbolic systems such as those governing fluid dynamics, errors can propagate by the convection-dominant nature of the system, making the prediction of regions requiring high resolution non-intuitive. An alternative and more robust way to identify the important areas for the output of interest is through adjoint solutions, which provide the sensitivity of the output to residual perturbations. The idea to combine output error estimation and gradient-based optimization is natural, as both methods require output adjoint solutions. Even though adjoint-based error estimation and mesh adaptation have been studied in depth and successfully

demonstrated in many aerospace engineering problems [15, 16, 17, 18], their application to optimization problems has received less attention. Lu [19] incorporated adjoint-based error estimation and order adaptation into gradient-based optimization. The constraints are realized as simple quadratic penalty functions added to the objective. Progressive optimization is used with mesh adaptation based on the error of the penalized objective. Nemeč and Aftosmis [20] modified the penalty terms to avoid vanishing of the constraint error when the constraints are satisfied. Li and Hartmann [21] eliminated the constraint by trimming with an individual design variable, and introduced a multi-target adaptation algorithm in which mesh adaptation targets the objective and constraint outputs equally on a fixed fidelity. Hicken and Alonso [22] used the gradient norm error as the refinement indicator, actively control the first-order optimality condition, while approximating higher-order derivatives. Chen and Fidkowski [23] adopted an efficient error estimate for the optimization problems accounting the effects of general output constraints, and used it to drive the progressive optimization.

In most previous works on optimization combined with error estimation, mesh adaptation has only been used to add refinement. However, in order to control the discretization error at each fidelity (error level), the mesh may get adapted in many areas that are important for the intermediate designs but that are not necessary for the final optimal design, which may decrease the efficiency of the high-fidelity optimization. Building on our previous work [23], we employ a more sophisticated spatial adaptation method: unstructured mesh optimization through error sampling and synthesis (MOESS) proposed by Yano and Darmofal [24]. A cost-based multi-level optimization framework is developed in this work. At each optimization level, the computational cost is restricted to a certain number of degrees of freedom. Within the same optimization fidelity, the mesh is optimized for each design to predict both the objective and constraint outputs accurately. The goal is to make the best use of a given number of degrees of freedom, so that the overall computational cost of the high-fidelity optimization can be reduced.

2. Optimization formulation

2.1. Continuous and discrete optimization

An aerodynamic shape optimization problem can be stated as a search for the design \mathbf{x} over the design space \mathcal{X} that minimizes a given objective function \mathcal{J} ,

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathcal{J}(\mathbf{u}, \mathbf{x}), \quad \mathbf{u} \in \mathcal{U}, \mathbf{x} \in \mathcal{X} \\ \text{s.t.} \quad & \mathcal{R}^e(\mathbf{u}, \mathbf{x}) = \mathbf{0} \\ & \mathcal{R}^{ie}(\mathbf{u}, \mathbf{x}) \geq \mathbf{0} \end{aligned} \quad (1)$$

where $\mathcal{J}: \mathcal{U} \times \mathcal{X} \rightarrow \mathbb{R}$ represents a scalar objective function, $\mathcal{R}^e: \mathcal{U} \times \mathcal{X} \rightarrow \mathbb{R}^{n_e}$ and $\mathcal{R}^{ie}: \mathcal{U} \times \mathcal{X} \rightarrow \mathbb{R}^{n_{ie}}$ denote n_e equality and n_{ie} inequality constraints, respectively. The objective and constraints are always defined by the outputs (responses) of the flow equations, for example lift or drag, which consequently depend on the flow state variables \mathbf{u} . The state \mathbf{u} is the solution of the governing equations, lying in the solution space \mathcal{U} , which can be an infinite-dimensional space. In a variational setting, the governing equations, often the Euler or Navier-Stokes equations, can be represented by a semi-linear form,

$$\mathcal{R}(\mathbf{u}, \mathbf{v}; \mathbf{x}) = 0, \quad \forall \mathbf{v} \in \mathcal{V} \quad (2)$$

where \mathcal{V} denotes an appropriately-defined test space, and the semi-linear residual map $\mathcal{R}: \mathcal{U} \times \mathcal{V} \rightarrow \mathbb{R}$ corresponds to the weak formulation of the flow equations. The state $\mathbf{u} \in \mathcal{U}$ is solved within the design space \mathcal{X} to satisfy the governing equations, and this implicitly defines \mathbf{u} as a function of \mathbf{x} : $\mathbf{u} = \mathbf{u}(\mathbf{x})$. Moreover, the optimal design \mathbf{x} has to be in the feasible space $\mathcal{F}(\mathcal{X}) = \{\mathbf{x} \in \mathcal{X} : \mathcal{R}^e(\mathbf{u}, \mathbf{x}) = \mathbf{0}, \mathcal{R}^{ie}(\mathbf{u}, \mathbf{x}) \geq \mathbf{0}\}$ that satisfies the constraints. Depending on the optimization algorithm, intermediate designs in an iterative process may not be in the feasible set.

An adjoint solution $\boldsymbol{\psi} \in \mathcal{V}$ can be defined as the sensitivity of the output of interest to the residual perturbation caused by a state perturbation \mathbf{w} . The adjoint satisfies the following weak form [18],

$$\mathcal{J}'[\mathbf{u}](\mathbf{w}) + \mathcal{R}'[\mathbf{u}](\mathbf{w}, \boldsymbol{\psi}) = 0, \quad \forall \mathbf{w} \in \mathcal{U} \quad (3)$$

We generally cannot solve the PDEs in Eqn. (2) and Eqn. (3) analytically, and hence we discretize them on a finite-dimensional computational domain. Then the original flow and adjoint equations can be reformulated as: determine $\mathbf{u}_h \in \mathcal{U}_h$ and $\boldsymbol{\psi}_h \in \mathcal{V}_h$ such that

$$\mathcal{R}_h(\mathbf{u}_h, \mathbf{v}_h) = 0, \quad \forall \mathbf{v}_h \in \mathcal{V}_h \quad (4)$$

$$\mathcal{J}'_h[\mathbf{u}_h](\mathbf{w}_h) + \mathcal{R}'_h[\mathbf{u}_h](\mathbf{w}_h, \boldsymbol{\psi}_h) = 0, \quad \forall \mathbf{w}_h \in \mathcal{U}_h \quad (5)$$

where \mathcal{U}_h and \mathcal{V}_h are finite-dimensional functional spaces. The subscript h indicates a discretization of the equations, including the approximation order and the spatial mesh. We define an N -dimensional basis set assuming that the trial and test spaces are the same, $\mathcal{U}_h = \mathcal{V}_h = \text{span}\{\phi_i, i = 1, \dots, N\}$. The state and adjoint solutions can then be rewritten as linear combinations of the basis functions, $\mathbf{u}_h = \sum_{i=1}^N U_{h,i} \phi_i$, $\boldsymbol{\psi}_h = \sum_{i=1}^N \Psi_{h,i} \phi_i$, so that \mathbf{u}_h and $\boldsymbol{\psi}_h$ can be uniquely defined by their associated coefficient vectors $\mathbf{U}_h = [U_{h,i}]_{i=1}^N$ and $\boldsymbol{\Psi}_h = [\Psi_{h,i}]_{i=1}^N$. We can also define the residual vector as $\mathbf{R}_h = [\mathcal{R}_h(\mathbf{u}_h, \phi_i)]_{i=1}^N$ and the output as $J_h(\mathbf{U}_h) = \mathcal{J}_h(\mathbf{u}_h)$, so that the state and adjoint equations can be written as

$$\mathbf{R}_h(\mathbf{U}_h) = \mathbf{0} \quad (6)$$

$$\left(\frac{\partial J_h}{\partial \mathbf{U}_h} \right)^T + \left(\frac{\partial \mathbf{R}_h}{\partial \mathbf{U}_h} \right)^T \boldsymbol{\Psi}_h = \mathbf{0} \quad (7)$$

Consequently, a fully-discretized form of the original optimization problem, which augments the constraints with the state equations, can be stated as

$$\begin{aligned} \min_{\mathbf{x}} \quad & J_h(\mathbf{U}_h, \mathbf{x}), \quad \mathbf{U}_h \in \mathbb{R}^N, \quad \mathbf{x} \in \mathcal{X} \\ \text{s.t.} \quad & \mathbf{R}_h(\mathbf{U}_h, \mathbf{x}) = \mathbf{0} \\ & \mathbf{R}_h^e(\mathbf{U}_h, \mathbf{x}) = \mathbf{0} \\ & \mathbf{R}_h^{\text{ie}}(\mathbf{U}_h, \mathbf{x}) \geq \mathbf{0} \end{aligned} \quad (8)$$

The equality and inequality constraints are also vectorized in the equation above. In this paper, continuous and discrete optimization refer to the optimization governed by PDEs in continuous and discretized form, while in other contexts these terms may refer to the optimization with continuous and discrete design spaces.

2.2. Optimization via the adjoint

Inactive inequality constraints, $\mathbf{R}_{\text{ia}}^{\text{ie}}$, do not affect the optimization explicitly, while the active ones, $\mathbf{R}_{\text{a}}^{\text{ie}}$, behave like equality constraints. We omit the subscript h here for simpler exposition. In general, the inequality constraints can also be transformed into equality constraints with non-negative slack variables [25]. For simplicity, we only consider the active inequality constraints and the equality constraints, put together into one vector of dimension N_t as trim constraints, $(\mathbf{R}^{\text{trim}})^T = [(\mathbf{R}^e)^T \quad (\mathbf{R}_{\text{a}}^{\text{ie}})^T] \in \mathbb{R}^{N_t}$,

$$\mathbf{R}^{\text{trim}}(\mathbf{U}, \mathbf{x}) = \mathbf{J}^{\text{trim}}(\mathbf{U}, \mathbf{x}) - \bar{\mathbf{J}}^{\text{trim}} = \mathbf{0} \quad (9)$$

where $\bar{\mathbf{J}}^{\text{trim}}$ is a set of N_t target trim outputs, for example, the target lift in an airfoil drag minimization problem. In order to distinguish the trim outputs from the objective output, we denote the latter by J^{adapt} , as the objective output is directly targeted for adaptation.

The adjoint-based optimization is equivalent to searching for the stationary point of the Lagrangian function that augments the flow equations with trim constraints,

$$\mathcal{L}(\mathbf{U}, \mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = J^{\text{adapt}}(\mathbf{U}, \mathbf{x}) + \boldsymbol{\lambda}^T \mathbf{R}(\mathbf{U}, \mathbf{x}) + \boldsymbol{\mu}^T \mathbf{R}^{\text{trim}}(\mathbf{U}, \mathbf{x}) \quad (10)$$

where $\boldsymbol{\lambda} \in \mathbb{R}^N$ and $\boldsymbol{\mu} \in \mathbb{R}^{N_t}$ are the Lagrange multipliers associated with the state equations and the trim constraints, respectively.

By requiring stationarity with respect to the design variables \mathbf{x} , the states \mathbf{U} , and the Lagrange multipliers $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$, we arrive at the *first-order necessary condition for optimality*, or the *Karush-Kuhn-Tucker condition*,

$$\frac{\partial \mathcal{L}}{\partial \mathbf{x}} = \frac{\partial J^{\text{adapt}}}{\partial \mathbf{x}} + \boldsymbol{\lambda}^T \frac{\partial \mathbf{R}}{\partial \mathbf{x}} + \boldsymbol{\mu}^T \frac{\partial \mathbf{R}^{\text{trim}}}{\partial \mathbf{x}} = \mathbf{0} \quad (11a)$$

$$\frac{\partial \mathcal{L}}{\partial \mathbf{U}} = \frac{\partial J^{\text{adapt}}}{\partial \mathbf{U}} + \boldsymbol{\lambda}^T \frac{\partial \mathbf{R}}{\partial \mathbf{U}} + \boldsymbol{\mu}^T \frac{\partial \mathbf{R}^{\text{trim}}}{\partial \mathbf{U}} = \mathbf{0} \quad (11b)$$

$$\frac{\partial \mathcal{L}}{\partial \boldsymbol{\lambda}} = \mathbf{R}(\mathbf{U}, \mathbf{x}) = \mathbf{0} \quad (11c)$$

$$\frac{\partial \mathcal{L}}{\partial \boldsymbol{\mu}} = \mathbf{R}^{\text{trim}}(\mathbf{U}, \mathbf{x}) = \mathbf{0} \quad (11d)$$

In practice, the above *optimality condition* is typically not solved simultaneously, since the whole system can be large, coupled, nonlinear, and often ill-conditioned [26]. Instead, the four sub-systems are solved individually and then coupled for the optimization. Specifically, in our method, we solve the state equations for a given design each time; in other words, Eqn. (11c) is always satisfied during the optimization. Then we can choose $\boldsymbol{\lambda}$ such that Eqn. (11b) is enforced after each state solve,

$$\boldsymbol{\lambda}^T = - \left(\frac{\partial J^{\text{adapt}}}{\partial \mathbf{U}} + \boldsymbol{\mu}^T \frac{\partial \mathbf{R}^{\text{trim}}}{\partial \mathbf{U}} \right) \frac{\partial \mathbf{R}}{\partial \mathbf{U}}^{-1} = (\boldsymbol{\Psi}^{\text{adapt}} + \boldsymbol{\Psi}^{\text{trim}} \boldsymbol{\mu})^T \quad (12)$$

Eqn. (12) gives a coupled adjoint variable $\boldsymbol{\lambda}$ that incorporates the adjoints of both the objective and the trim outputs, $\boldsymbol{\Psi}^{\text{adapt}} \in \mathbb{R}^{N \times 1}$ and $\boldsymbol{\Psi}^{\text{trim}} \in \mathbb{R}^{N \times N_t}$, which satisfy the discretized adjoint equation in Eqn. (7)

$$\frac{\partial \mathbf{R}}{\partial \mathbf{U}}^T \boldsymbol{\Psi}^{\text{adapt}} + \frac{\partial J^{\text{adapt}}}{\partial \mathbf{U}}^T = \mathbf{0}, \quad \frac{\partial \mathbf{R}}{\partial \mathbf{U}}^T \boldsymbol{\Psi}^{\text{trim}} + \frac{\partial \mathbf{J}^{\text{trim}}}{\partial \mathbf{U}}^T = \mathbf{0} \quad (13)$$

With this specific choice of $\boldsymbol{\lambda}$, we can evaluate the gradient of the Lagrangian function with respect to the design variables, starting with Eqn. (11a),

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \mathbf{x}} &= \frac{\partial J^{\text{adapt}}}{\partial \mathbf{x}} + \boldsymbol{\lambda}^T \frac{\partial \mathbf{R}}{\partial \mathbf{x}} + \boldsymbol{\mu}^T \frac{\partial \mathbf{R}^{\text{trim}}}{\partial \mathbf{x}} \\ &= \frac{\partial J^{\text{adapt}}}{\partial \mathbf{x}} + (\boldsymbol{\Psi}^{\text{adapt}})^T \frac{\partial \mathbf{R}}{\partial \mathbf{x}} + \boldsymbol{\mu}^T \left[\frac{\partial \mathbf{R}^{\text{trim}}}{\partial \mathbf{x}} + (\boldsymbol{\Psi}^{\text{trim}})^T \frac{\partial \mathbf{R}}{\partial \mathbf{x}} \right] \\ &= \frac{dJ^{\text{adapt}}}{d\mathbf{x}} + \boldsymbol{\mu}^T \frac{d\mathbf{J}^{\text{trim}}}{d\mathbf{x}} \end{aligned} \quad (14)$$

The last equality is obtained via adjoint-based sensitivity analysis, where $d(\cdot)/d\mathbf{x}$ denotes the total derivative with respect to design variables by considering the states as an implicit function of \mathbf{x} , $\mathbf{U}(\mathbf{x})$.

Now the optimization problem has been reduced to finding an optimal design \mathbf{x} and the corresponding Lagrange multipliers $\boldsymbol{\mu}$ satisfying,

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \mathbf{x}} &= \frac{dJ^{\text{adapt}}}{d\mathbf{x}} + \boldsymbol{\mu}^T \frac{d\mathbf{J}^{\text{trim}}}{d\mathbf{x}} = \mathbf{0} \\ \frac{\partial \mathcal{L}}{\partial \boldsymbol{\mu}} &= \mathbf{R}^{\text{trim}} = \mathbf{0} \end{aligned} \quad (15)$$

However, since in a practical calculation, on a finite-dimensional space, the discretization error appears in both the flow equations and the adjoint equations, optimality cannot be guaranteed even when Eqn. (15) is satisfied. The present work focuses on controlling the error in the optimization problem via error estimation and mesh adaptation.

3. Governing equations and discretization

Evaluation of the objective and the constraints at each optimization step relies on the numerical simulation of the system. The governing PDE considered here is transport in conservation form

$$\frac{\partial \mathbf{u}}{\partial t} + \nabla \cdot \vec{\mathbf{F}}(\mathbf{u}, \nabla \mathbf{u}) + \mathbf{S}(\mathbf{u}, \nabla \mathbf{u}) = \mathbf{0} \quad (16)$$

which encompasses scalar advection-diffusion and the compressible Navier-Stokes equations. For the latter, \mathbf{u} is the conservative state vector composed by the flow variables, $\vec{\mathbf{F}}$ denotes the total inviscid and viscous flux vectors, and \mathbf{S} represents the source term required when modeling turbulence. When running Reynolds-averaged turbulent cases, we use the Spalart-Allmaras one-equation model, with a negative turbulent-viscosity modification [27].

We discretize Eqn. (16) with the discontinuous Galerkin (DG) finite-element method, which is suitable for high-order accuracy and hp -refinement [28, 29, 30]. The computational domain Ω is divided into a shape-regular mesh \mathcal{T}_h consisting of N_e non-overlapping elements Ω_e , $\mathcal{T}_h = \{\Omega_e : \bigcup_{e=1}^{N_e} \Omega_e = \Omega, \bigcap_{e=1}^{N_e} \Omega_e = \emptyset\}$. In DG, the state is approximated by piece-wise polynomials lying on the approximation space \mathcal{V}_h^p , with no continuity constraints imposed on the approximations between adjacent elements. The approximation space consists of element-wise polynomials and is defined as $\mathcal{V}_h^p = \{v_h \in L^2(\Omega) : v_h|_{\Omega_e} \in \mathcal{P}^{p_e}, \forall \Omega_e \in \mathcal{T}_h\}$, where \mathcal{P}^{p_e} denotes polynomials of order p_e on element Ω_e , a distribution that is not necessary uniform throughout the mesh. The weak form of Eqn. (16) follows from multiplying the equation by test functions (taken from the approximation space), integrating by parts, and coupling elements via unique inter-element fluxes. We use the Roe approximate Riemann solver [31] for the inviscid flux, and the second form of Bassi and Rebay (BR2) for viscous flux [32]. Choosing a basis for the test and trial spaces yields a system of nonlinear, algebraic equations in the form of Eqn. (6)

$$\mathbf{R}_H(\mathbf{U}_H, \mathbf{x}) = \mathbf{0} \quad (17)$$

Here, \mathbf{R}_H is the residual vector, a nonlinear function of the discrete state vector \mathbf{U}_H and the design variables \mathbf{x} . For the steady state problems considered in this work, \mathbf{R}_H is the discrete spatial residual vector. The subscript H refers to discretization fidelity of the approximation/test space with respect to the approximation order and mesh refinement.

4. Objective error estimation

4.1. Adjoint-based error estimation

In practice it is not possible to obtain the true discretization error for an output, whereas the difference between a coarse space and fine space solution serves as an acceptable surrogate,

$$\text{output error: } \delta J \equiv J_H(\mathbf{U}_H) - J_h(\mathbf{U}_h) \quad (18)$$

In this expression, J represents the output of interest, and the subscripts h and H denote the fine and coarse spaces, respectively. In the present work, the fine space is achieved by increasing the elements' approximation order p_e , to $p_e + 1$. We do not solve the nonlinear fine-space flow problem for the error prediction, and instead we use the linear fine-space adjoint solution, Ψ_h . The adjoint weights the residual perturbation to produce an output perturbation,

$$\begin{aligned} \delta J &= J_H(\mathbf{U}_H) - J_h(\mathbf{U}_h) \\ &= J_h(\mathbf{U}_h^H) - J_h(\mathbf{U}_h) = \frac{\partial J_h}{\partial \mathbf{U}_h} \delta \mathbf{U} \\ &= -\Psi_h^T \delta \mathbf{R}_h = -\Psi_h^T [\mathbf{R}_h(\mathbf{U}_h^H) - \mathbf{R}_h(\mathbf{U}_h)] \\ &= -\Psi_h^T \mathbf{R}_h(\mathbf{U}_h^H) \end{aligned} \quad (19)$$

where \mathbf{U}_h is the (hypothetical) exact solution on the fine space, and \mathbf{U}_h^H is the state injected into the fine space from the coarse one, which generally will not give a zero fine space residual, $\mathbf{R}_h(\mathbf{U}_h^H) \neq \mathbf{R}_h(\mathbf{U}_h) = \mathbf{0}$. The derivation of Eqn. (19) originates from the small perturbation assumptions, and is valid for outputs whose definition does not change between the coarse and fine spaces, $J_H(\mathbf{U}_H) = J_h(\mathbf{U}_h^H)$.

4.2. Error estimation for optimization problems

Normally, error estimation is applied only to the output in which we are most interested, i.e. the objective. However, our optimization problem requires the simultaneous solution of flow equations and constraints, i.e. trim outputs. The discretization error of the trim outputs may indirectly affect the calculation of the objective [33]. To take this effect into account, the coupled adjoint should be used for the error estimation.

Consider a given design $\mathbf{x} = \mathbf{x}_H$, and suppose that the error of the objective only comes from the inexact solution \mathbf{U}_h^H . We can estimate the error in the objective with the linearization given by Eqn. (11b),

$$\begin{aligned}
\delta J^{\text{adapt}}(\mathbf{x}_H) &= J_h^{\text{adapt}}(\mathbf{U}_h^H, \mathbf{x}_H) - J_h^{\text{adapt}}(\mathbf{U}_h, \mathbf{x}_H) = \frac{\partial J_h^{\text{adapt}}}{\partial \mathbf{U}_h} \delta \mathbf{U} \\
&= -\boldsymbol{\lambda}_h^T \frac{\partial \mathbf{R}_h}{\partial \mathbf{U}_h} \delta \mathbf{U} - \boldsymbol{\mu}_h^T \frac{\partial \mathbf{R}_h^{\text{trim}}}{\partial \mathbf{U}_h} \delta \mathbf{U} = -\boldsymbol{\lambda}_h^T \delta \mathbf{R}_h - \boldsymbol{\mu}_h^T \delta \mathbf{R}_h^{\text{trim}} \\
&= -\boldsymbol{\lambda}_h^T [\mathbf{R}_h(\mathbf{U}_h^H, \mathbf{x}_H) - \mathbf{R}_h(\mathbf{U}_h, \mathbf{x}_H)] - \boldsymbol{\mu}_h^T [\mathbf{R}_h^{\text{trim}}(\mathbf{U}_h^H, \mathbf{x}_H) - \mathbf{R}_h^{\text{trim}}(\mathbf{U}_h, \mathbf{x}_H)] \\
&= -\boldsymbol{\lambda}_h^T \mathbf{R}_h(\mathbf{U}_h^H, \mathbf{x}_H) - \boldsymbol{\mu}_h^T [\mathbf{J}_h^{\text{trim}}(\mathbf{U}_h^H, \mathbf{x}_H) - \mathbf{J}_h^{\text{trim}}(\mathbf{U}_h, \mathbf{x}_H)] \\
&= -\left(\boldsymbol{\Psi}_h^{\text{adapt}} + \boldsymbol{\Psi}_h^{\text{trim}} \boldsymbol{\mu}_h \right)^T \mathbf{R}_h(\mathbf{U}_h^H, \mathbf{x}_H) - \boldsymbol{\mu}_h^T \delta \mathbf{J}^{\text{trim}}(\mathbf{x}_H) \\
&= -\left(\boldsymbol{\Psi}_h^{\text{adapt}} \right)^T \mathbf{R}_h(\mathbf{U}_h^H, \mathbf{x}_H) - \boldsymbol{\mu}_h^T \left(\boldsymbol{\Psi}_h^{\text{trim}} \right)^T \mathbf{R}_h(\mathbf{U}_h^H, \mathbf{x}_H) - \boldsymbol{\mu}_h^T \delta \mathbf{J}^{\text{trim}}(\mathbf{x}_H) \\
&= -\left(\boldsymbol{\Psi}_h^{\text{adapt}} \right)^T \mathbf{R}_h(\mathbf{U}_h^H, \mathbf{x}_H) + \boldsymbol{\mu}_h^T \delta \mathbf{J}^{\text{trim}}(\mathbf{x}_H) - \boldsymbol{\mu}_h^T \delta \mathbf{J}^{\text{trim}}(\mathbf{x}_H) \\
&= -\left(\boldsymbol{\Psi}_h^{\text{adapt}} \right)^T \mathbf{R}_h(\mathbf{U}_h^H, \mathbf{x}_H)
\end{aligned} \tag{20}$$

where $\delta \mathbf{J}^{\text{trim}}(\mathbf{x}_H)$ is a vector containing error estimates of the N_t constraint outputs at the fixed design \mathbf{x}_H . Following the definition of the output error in Eqn. (19), $\delta \mathbf{J}^{\text{trim}}(\mathbf{x}_H) = \mathbf{J}_H^{\text{trim}}(\mathbf{U}_H, \mathbf{x}_H) - \mathbf{J}_h^{\text{trim}}(\mathbf{U}_h, \mathbf{x}_H) = \mathbf{J}_h^{\text{trim}}(\mathbf{U}_h^H, \mathbf{x}_H) - \mathbf{J}_h^{\text{trim}}(\mathbf{U}_h, \mathbf{x}_H) = -\left(\boldsymbol{\Psi}_h^{\text{trim}} \right)^T \mathbf{R}_h(\mathbf{U}_h^H, \mathbf{x}_H)$. Eqn. (20) is consistent with the previous analysis without the trim conditions, since we keep the design fixed between the coarse and fine spaces, and because we assume that the error only comes from the inexact state solution \mathbf{U}_h^H . In general, however, we need to deal with both the objective error and the constraints error. The problem becomes worse if we have high accuracy in the objective while little confidence in the constraint outputs, or vice versa. If we run the optimization on the fine space and the coarse space, even with the same target constraint outputs, we will generally obtain different designs. This difference may come from the deviation of both the design parameters and the flow states, and separate error estimation and mesh adaptation for the objective and trim outputs can be inefficient.

If we consider the optimal design on the coarse space $(\mathbf{U}_H, \mathbf{x}_H)$ and the fine space $(\mathbf{U}_h, \mathbf{x}_h)$, since the optimality conditions Eqn. (11a) and Eqn. (11b) both hold now, the error from the inexact solution \mathbf{U}_h^H as

well as the deficient design \mathbf{x}_H can be estimated as

$$\begin{aligned}
\delta J_{\text{opt}}^{\text{adapt}} &= J_h^{\text{adapt}}(\mathbf{U}_h^H, \mathbf{x}_H) - J_h^{\text{adapt}}(\mathbf{U}_h, \mathbf{x}_h) = \frac{\partial J_h^{\text{adapt}}}{\partial \mathbf{U}_h} \delta \mathbf{U} + \frac{\partial J_h^{\text{adapt}}}{\partial \mathbf{x}} \delta \mathbf{x} \\
&= -\boldsymbol{\lambda}_h^T \left(\frac{\partial \mathbf{R}_h}{\partial \mathbf{U}_h} \delta \mathbf{U} + \frac{\partial \mathbf{R}_h}{\partial \mathbf{x}} \delta \mathbf{x} \right) - \boldsymbol{\mu}_h^T \left(\frac{\partial \mathbf{R}_h^{\text{trim}}}{\partial \mathbf{U}_h} \delta \mathbf{U} + \frac{\partial \mathbf{R}_h^{\text{trim}}}{\partial \mathbf{x}} \delta \mathbf{x} \right) = -\boldsymbol{\lambda}_h^T \delta \mathbf{R}_h - \boldsymbol{\mu}_h^T \delta \mathbf{R}_h^{\text{trim}} \\
&= -\boldsymbol{\lambda}_h^T [\mathbf{R}_h(\mathbf{U}_h^H, \mathbf{x}_H) - \mathbf{R}_h(\mathbf{U}_h, \mathbf{x}_h)] - \boldsymbol{\mu}_h^T [\mathbf{R}_h^{\text{trim}}(\mathbf{U}_h^H, \mathbf{x}_H) - \mathbf{R}_h^{\text{trim}}(\mathbf{U}_h, \mathbf{x}_h)] \\
&= -\boldsymbol{\lambda}_h^T \mathbf{R}_h(\mathbf{U}_h^H, \mathbf{x}_H) - \boldsymbol{\mu}_h^T \mathbf{R}_h^{\text{trim}}(\mathbf{U}_h^H, \mathbf{x}_H) \\
&= -\left(\boldsymbol{\Psi}_h^{\text{adapt}} + \boldsymbol{\Psi}_h^{\text{trim}} \boldsymbol{\mu}_h \right)^T \mathbf{R}_h(\mathbf{U}_h^H, \mathbf{x}_H) - \boldsymbol{\mu}_h^T [\mathbf{J}_h^{\text{trim}}(\mathbf{U}_h^H, \mathbf{x}_H) - \bar{\mathbf{J}}^{\text{trim}}] \\
&= -\left(\boldsymbol{\Psi}_h^{\text{adapt}} \right)^T \mathbf{R}_h(\mathbf{U}_h^H, \mathbf{x}_H) - \boldsymbol{\mu}_h^T \left(\boldsymbol{\Psi}_h^{\text{trim}} \right)^T \mathbf{R}_h(\mathbf{U}_h^H, \mathbf{x}_H) - \boldsymbol{\mu}_h^T [\mathbf{J}_h^{\text{trim}}(\mathbf{U}_h^H, \mathbf{x}_H) - \bar{\mathbf{J}}^{\text{trim}}] \\
&= \delta J^{\text{adapt}}(\mathbf{x}_H) + \boldsymbol{\mu}_h^T \delta \mathbf{J}^{\text{trim}}(\mathbf{x}_H) - \boldsymbol{\mu}_h^T [\mathbf{J}_h^{\text{trim}}(\mathbf{U}_h^H, \mathbf{x}_H) - \bar{\mathbf{J}}^{\text{trim}}]
\end{aligned} \tag{21}$$

Since the trim constraints are satisfied for the optimal design on the coarse and fine spaces respectively, i.e. $(\mathbf{U}_H, \mathbf{x}_H)$ and $(\mathbf{U}_h, \mathbf{x}_h)$, and since the definition of the outputs is often the same on the coarse and fine spaces, we generally have that

$$\mathbf{J}_h^{\text{trim}}(\mathbf{U}_h^H, \mathbf{x}_H) = \mathbf{J}_H^{\text{trim}}(\mathbf{U}_H, \mathbf{x}_H) = \bar{\mathbf{J}}^{\text{trim}} = \mathbf{J}_h^{\text{trim}}(\mathbf{U}_h, \mathbf{x}_h) \tag{22}$$

Hence, the last term in Eqn. (21) is often negligible for the optimal design, resulting a simpler form for the error of the optimal objective,

$$\delta J_{\text{opt}}^{\text{adapt}} = -\left(\boldsymbol{\Psi}_h^{\text{adapt}} + \boldsymbol{\Psi}_h^{\text{trim}} \boldsymbol{\mu}_h \right)^T \mathbf{R}_h(\mathbf{U}_h^H, \mathbf{x}_H) = \delta J^{\text{adapt}}(\mathbf{x}_H) + \boldsymbol{\mu}_h^T \delta \mathbf{J}^{\text{trim}}(\mathbf{x}_H) \tag{23}$$

The error without the subscript ‘‘opt’’ is the output error without the trim constraints. Eqn. (23) gives a prediction of the objective output error due to the finite spatial discretization, for calculations on the inaccurate optimal design using the coarse discretization space. With the advantages of adjoint-based error estimation, we avoid the expensive solves of both the optimal design \mathbf{x}_h and flow states \mathbf{U}_h , i.e. the whole optimization process on the fine space. This is because the error estimate requires only residuals, which can be computed from \mathbf{x}_H and \mathbf{U}_H , without the fine-space flow states or the optimal design. However, the estimation requires the fine-space adjoints $\boldsymbol{\Psi}_h$ as well as the fine-space Lagrange multipliers $\boldsymbol{\mu}_h$. In our implementation, the fine-space adjoints $\boldsymbol{\Psi}_h$ are approximated by reconstructing the coarse-space adjoints $\boldsymbol{\Psi}_H$ [18], while the Lagrange multipliers are extracted from the optimizer on the coarse space.

Eqn. (23) provides the objective error estimate of the optimal design, which does not always hold during the optimization process. Thus, it is neither the error of the objective nor the error of the constraints when the design is away from optimal. However, it couples the objective error and constraints error, giving a better error level for the whole optimization problem, so it is expected to serve as a better adaptation indicator for the optimization or constrained problems. Moreover, in the multifidelity optimization framework, when most of the mesh adaptation happens after a successful optimization on the current fidelity, using Eqn. (23) can be more efficient than using only the objective error without constraints [33]. In this paper, we allow violation of the constraints during the optimization process. The constraints are only enforced for the optimal design, while for the methods requiring the exploration path to be always feasible, Eqn. (22) holds for every iteration, thereby the objective error estimate in Eqn. (23) is valid for all the intermediate designs.

5. Mesh refinement and mesh optimization

5.1. Adaptation indicator

For a single output of interest, the error estimate can be localized in each element and serves as an indicator for mesh adaptation. A common approach is to keep track of the elemental error contribution,

taking its absolute value as the indicator.

$$\mathcal{E} = J_H(\mathbf{U}_H) - J_h(\mathbf{U}_h) = -\Psi_h^T \mathbf{R}_h(\mathbf{U}_h^H) = -\sum_{e=1}^{N_e} \Psi_{h,e}^T \mathbf{R}_{h,e}(\mathbf{U}_h^H), \quad \epsilon_e \equiv \left| \Psi_{h,e}^T \mathbf{R}_{h,e}(\mathbf{U}_h^H) \right| \quad (24)$$

where \mathcal{E} denotes the total output error estimate, $\epsilon_e \geq 0$ is the error indicator for element Ω_e .

For the optimization problem, the error estimate in Eqn. (23) can also be localized in each element and guides the mesh adaptation.

$$\epsilon_e = \left| -(\Psi_{h,e}^{\text{adapt}} + \Psi_{h,e}^{\text{trim}} \boldsymbol{\mu}_h)^T \mathbf{R}_{h,e}(\mathbf{U}_h^H, \mathbf{x}_H) \right| = \left| \delta J_{h,e}^{\text{adapt}} + \boldsymbol{\mu}_h^T \delta \mathbf{J}_{h,e}^{\text{trim}} \right| \leq \epsilon_e^{\text{adapt}} + |\boldsymbol{\mu}_h^T| \epsilon_e^{\text{trim}} = \epsilon_{e,\text{con}} \quad (25)$$

where ϵ_e allows cancellations between objective and constraints error indicators, while $\epsilon_{e,\text{con}}$ provides a more conservative error indicator.

A naive adaptation strategy is to adapt a fixed fraction of the mesh at each adaptation cycle, in which the elements with highest errors are flagged for refinement. However, such a strategy cannot detect strong directional features, such as shocks or boundary layers in flow problems. In order to provide anisotropic resolution, we have to include the solution anisotropy with the error estimate in the adaptation techniques. We consider two anisotropic adaptation strategies here: 1) goal-oriented mesh adaptation with Hessian-based anisotropy detection, and 2) mesh optimization via error sampling and synthesis (MOESS).

5.2. Metric-based remeshing

A Riemannian metric field, $\mathcal{M}(\vec{x})$, is a smoothly varying field of symmetric positive definite (SPD) tensors that can be used to encode anisotropic information of the computational mesh, including desired mesh sizes and stretching directions. At any point in the physical space, \vec{x} , the metric tensor $\mathcal{M}(\vec{x})$ provides a ‘‘yardstick’’ for measuring the distance from \vec{x} to another point infinitesimally far away, $\vec{x} + \delta\vec{x}$. The distance under the Riemannian metric is given by

$$\delta l = \sqrt{\delta\vec{x}^T \mathcal{M}(\vec{x}) \delta\vec{x}} \quad (26)$$

After choosing a Cartesian coordinate system and basis for d -dimensional physical space, \mathcal{M} can be represented as a $d \times d$ SPD matrix. The set of points at unit metric distance is an ellipse in 2D or an ellipsoid in 3D: eigenvectors of \mathcal{M} gives its principal axes, while the length of each axis (stretching) is the inverse square root of the corresponding eigenvalue.

A mesh that conforms to a metric field is one in which all the edges have unit length under the metric, to some tolerance. The metric-conforming mesh is not unique; however, a family of *metric-conforming* meshes have similar approximation properties [34]. The *metric-conforming* mesh generator used in this work is the Bi-dimensional Anisotropic Mesh Generator (BAMG) [35]. BAMG requires a metric field which is specified at vertices of a background mesh to generate a new mesh described by the continuous metric field. For a desired new d -dimensional simplex mesh, the *mesh-implied* metric can be obtained by solving a linear system for the $d(d+1)/2$ independent entries of \mathcal{M}_e at each element. The equations in this system enforce that each of the $d(d+1)/2$ edges has unit metric length. Then the discontinuous elemental metric field is averaged to the surrounding vertices using an affine-invariant algorithm [36]. The metric-conforming mesh and mesh-implied metric give a way of converting between an anisotropic mesh and a Riemannian metric field, a so called mesh-metric duality.

5.3. Goal-oriented mesh adaptation with Hessian-based anisotropy detection

Given a localized error estimate, an appropriate adaptation strategy can be determined by decreasing and equally distributing the error [10]. In order to make the metric-based mesh adaptation efficient, stretched elements have to be generated in areas where the solution exhibits high anisotropy. One dominant method for detecting the anisotropy is to estimate directional interpolation error of a scalar solution [10, 37, 38]. For

approximation order $p = 1$, the error of a scalar solution u over an edge E in the mesh, with unit tangent vector \mathbf{s} and length h , is given by

$$\epsilon_E \propto |\mathbf{s}^T \mathbf{H} \mathbf{s}| h^2 \quad (27)$$

where \mathbf{H} is the solution Hessian matrix,

$$H_{i,j} = \frac{\partial^2 u}{\partial x_i \partial x_j}, \quad i, j \in [1, \dots, d] \quad (28)$$

Suppose the edge conforms to a metric \mathcal{M} , assumed constant along the edge, then the edge is of unit length under the metric measure,

$$l_{\mathcal{M}} = \sqrt{\mathbf{s}^T \mathcal{M} \mathbf{s}} h = 1 \quad (29)$$

Thus the interpolation error and metric are related by Eqn. (27) and Eqn. (29), and with the requirements of error equidistribution, we have

$$\frac{|\mathbf{s}^T \mathbf{H} \mathbf{s}|}{\mathbf{s}^T \mathcal{M} \mathbf{s}} = C \quad (30)$$

where C is a constant defined by the desired error distribution. In order for Eqn. (30) to be valid for edges in any principal direction, \mathcal{M} can be chosen as

$$\mathcal{M}_{\mathbf{H}} = \frac{1}{C} \mathbf{Q} |\mathbf{\Lambda}| \mathbf{Q}^T = \frac{1}{h_{\text{ref}}^2} \mathbf{Q} |\mathbf{\Lambda}| \mathbf{Q}^T \quad (31)$$

Here, \mathbf{Q} denotes the orthonormal matrix containing the eigenvectors of \mathbf{H} , and $\mathbf{\Lambda}$ as the corresponding diagonal matrix containing its eigenvalues. h_{ref} controls the absolute mesh size, which can be determined by the error estimates and desired error distribution. Many previous works have studied Hessian-based anisotropic mesh adaptation [34, 39], in which the local mesh size or mesh density is determined by equally distributing the solution interpolation error. In this work, we adopt the approach developed in [38], where h_{ref} is obtained by output error equidistribution. In order to distinguish the current approach from pure Hessian-based anisotropic mesh adaptation, we refer to this approach as goal-oriented mesh adaptation with Hessian-based anisotropy detection, or goal-oriented Hessian-based anisotropic mesh adaptation. Anisotropy detection based on the standard Hessian matrix is not suited for higher order interpolation, due to the linear interpolation assumption used in the derivation of the Hessian-based method. Fidkowski and Darmofal [40] extended the Hessian-based anisotropy detection to general approximation order p by estimating the $p + 1^{\text{st}}$ derivatives.

In order to equally distribute the output error, we also need to predict the element size, or the number of the elements N^f , in the adapted (fine) mesh. Let n_k , not necessarily an integer, be the number of the fine-mesh elements contained in element k at the original mesh. Denoting the current element size by h_{ref}^c and the requested element size as h_{ref}^f , n_k can be approximated as

$$n_k = \left(\frac{h_{\text{ref}}^c}{h_{\text{ref}}^f} \right)^d \quad (32)$$

Given an output error tolerance e_0 , to satisfy the error equidistribution, each fine-mesh element is allowed an error of e_0/N^f , which means that each element k is allowed an error of $n_k e_0/N^f$. We relate the growth in elements to an error reduction factor through an a priori estimate

$$\underbrace{n_k \frac{e_0}{N^f}}_{\text{allowable error}} = \underbrace{\epsilon_k \left(\frac{h_{\text{ref}}^f}{h_{\text{ref}}^c} \right)^{\bar{p}_k + 1}}_{\text{a priori estimate}} \quad (33)$$

where ϵ_k is the current error indicator, $\bar{p}_k = \min(p_k, \gamma_k)$, and γ_k is the lowest order of any singularity within element k . Substituting Eqn. (32) into Eqn. (33) yields a relation between n_k and N^f .

$$n_k \frac{e_0}{N^f} = \epsilon_k n_k^{-(\bar{p}_k + 1)/d} \Rightarrow n_k^{1 + (\bar{p}_k + 1)/d} = \frac{\epsilon_k}{e_0/N^f} \quad (34)$$

Substituting Eqn. (34) into $N^f = \sum_k n_k$, we can solve for N^f . In practice, we use a fixed-growth refinement strategy instead of relying on the a priori error estimate, i.e., assume $N^f = f^{\text{growth}} N^c$ at each adaptation iteration, n_k and h_{ref}^f are determined by Eqn. (34) and Eqn. (32) respectively. Adaptation stops when we meet the error tolerance, $\epsilon \equiv \sum_k \epsilon_k \leq \epsilon_0$. For high-order solutions, the first d principal axes of the $p + 1$ directional derivatives are used to characterize the anisotropy and to guide the goal-oriented mesh adaptation.

5.4. Mesh optimization through error sampling and synthesis (MOESS)

Goal-oriented Hessian-based anisotropic mesh adaptation has been shown to successfully detect solution anisotropy in many applications. However, it relies on a scalar solution u , which should be carefully chosen to correlate to the chosen output of interest. Also, an inflection in u may lead to inappropriate mesh stretching, and inadequate resolution may occur where the magnitude of the Hessian is close to zero. In addition, the fixed-fraction adaptation strategy and the stopping criterion used in goal-oriented Hessian-based adaptation can cause over-refinement for a design evaluation with certain error requirements. Due to these deficiencies, extra degrees of freedom may need to be added when the design changes. This kind of cost allocation may accumulate during the optimization, resulting in inefficiencies in practical applications.

In order to maximize the approximation potential of a mesh with a given number of degrees of freedom, we consider a more sophisticated spatial adaptation method: unstructured mesh optimization through error sampling and synthesis (MOESS). In MOESS, the mesh adaptation is formulated as an optimization problem in which the optimal change of the metric field is iteratively determined based on a prescribed metric-cost model and a sampling-inferred metric-error relationship. We briefly review this method and discuss its modifications in this section.

5.4.1. Error convergence model

The mesh optimization algorithm requires a model for how the output error changes as the metric changes. Suppose that a metric step matrix \mathcal{S}_e is imposed on an element Ω_e with current error indicator of ϵ_{e0} . Instead of using a typical priori model in Eqn. (33), a generalized error model taking account of the output's directional convergence property is given in [24],

$$\epsilon_e = \epsilon_{e0} \exp[\text{tr}(\mathcal{R}_e \mathcal{S}_e)] \quad \Rightarrow \quad \frac{\partial \epsilon_e}{\partial \mathcal{S}_e} = \epsilon_e \mathcal{R}_e \quad (35)$$

where \mathcal{R}_e is a symmetric error convergence rate tensor containing the directional convergence information, while the step matrix \mathcal{S}_e encodes both the size change and the stretching of the new element. The total output error over the mesh is the sum of the elemental errors, $\epsilon = \sum_{e=1}^{N_e} \epsilon_e$. During the optimization we will want to keep ϵ small, and we will want to determine the optimal step matrices at each element, and hence the metric changes at mesh vertices, \mathcal{S}_v . The rate tensor, \mathcal{R}_e , is determined separately for each element through a local output error sampling procedure in which the element is refined in different configurations and the resulting changes to the output error are estimated.

For a triangular element, we consider four refinement options, indexed by i , as shown in Figure 2. We would like to know how much the error would decrease under each refinement option. One expensive option is to refine the element with the proposed cut, re-solve the primal and fine-space adjoint problems globally, and re-compute the error estimate. Though accurate, this would be impractically expensive. Another option is to only solve the primal/adjoint problems on a subset of the original mesh: the current element and its neighbors. This approach, taken in [24], is less accurate but still performs very well as globally-exact primal/adjoint states are not necessary to estimate the error rate tensor. In this work we further simplify the estimation by not solving additional problems, even on a local patch of elements. Instead, we use an element-local projection method [41] to approximate the fine-space adjoint in semi-refined spaces associated with each refinement option.




Figure 2: Four refinement options for a triangle. Each one is considered implicitly during error sampling, though the elements are never actually refined.

5.4.2. Cost model

To measure the cost of refinement, we use degrees of freedom, dof, which on each element just depends on the approximation order p_e , assumed constant and equal to p over the elements. Again, we consider one element, Ω_e , with current cost C_{e0} , and a proposed metric step matrix \mathcal{S}_e . The cost allocation for the new configuration is inversely proportional to the element area, which can be inferred from the metric,

$$C_e = C_{e0} \sqrt{\frac{\det(\mathcal{M}_0^{1/2} \exp(\mathcal{S}_e) \mathcal{M}_0^{1/2})}{\det(\mathcal{M}_0)}} = C_{e0} \exp\left[\frac{1}{2} \text{tr}(\mathcal{S}_e)\right] \Rightarrow \frac{\partial C_e}{\partial \mathcal{S}_e} = C_e \frac{1}{2} \mathcal{I} \quad (36)$$

where C_e is the expected cost over the original element area after applying \mathcal{S}_e to the original metric. The total cost over the mesh is the sum of the elemental costs, $C = \sum_{e=1}^{N_e} C_e$. During the optimization, we will want to determine the optimal step matrices given a fixed cost to minimize the total error indicator ϵ .

5.4.3. Mesh optimization algorithm

Given a current mesh with its mesh-implied metric ($\mathcal{M}_0(\vec{x})$), elemental error indicator ϵ_{e0} , and the elemental rate tensor \mathcal{R}_e , the mesh optimization problem can be formulated as

$$\begin{aligned} \min_{\mathcal{S}_v} \quad & \epsilon(\mathcal{S}_v) \\ \text{s.t.} \quad & C(\mathcal{S}_v) = \text{const} \end{aligned} \quad (37)$$

where \mathcal{S}_v is the step matrix at each vertex, and the elemental step matrix \mathcal{S}_e in the error and cost models takes the algorithmic mean of its vertices' step matrices. Again, using the Lagrangian function, we have the first order optimality condition as

$$\frac{\partial \epsilon}{\partial \mathcal{S}_v} + \lambda_s \frac{\partial C}{\partial \mathcal{S}_v} = \mathbf{0} \quad (38)$$

where λ_s is the global Lagrange multiplier, taking the same value in every element. We do not solve this problem exactly, since this would be a very high-dimensional problem which may require extremely high computational effort, especially in an optimization problem where the optimal mesh changes as the design varies. Furthermore, the error model based on the empirical local sampling may not represent the error exactly. Therefore, solving the mesh optimization problem exactly is inefficient and unnecessary. Instead, we follow Yano's optimization approach, by defining a "local" Lagrange multiplier,

$$\lambda_v = \frac{\partial \epsilon / \partial s_v}{\partial C / \partial s_v} \quad (39)$$

where $s_v = \text{tr}(\mathcal{S}_v)$ is the trace of the step matrix. λ_v can be interpreted as the marginal improvement in the local error for a given investment in the local cost. The optimization problem is to equally distribute λ_v over the mesh vertices, which eventually converges to the same solution as Eqn. (38). In practice, the mesh optimization and state/adjoint solution are performed several times at a given target cost, until the error stops changing. Then the target cost is increased to reduce the error further if desired.

6. Optimization approach

6.1. Optimization algorithm

To be consistent with the analysis for the optimization formulation in Section 2, the optimization algorithm should involve the Lagrange multipliers. Sequential Least Square Quadratic Programming (SLSQP) [42] with Broyden-Fletcher-Goldfarb-Shanno (BFGS) type Hessian approximation [43] is used in this work. The weak Wolfe condition is used to terminate the backtracking line search, ensuring a sufficient decrease at each optimization step. The gradients of the objective function and the constraints are calculated by the adjoint method, per Eqn. (14), and the objective and constraints are evaluated with the numerical solution of the state problem, Eqn. (16). The Lagrange multipliers associated with the trim constraints are extracted after each optimization step as a surrogate for the fine space multipliers, used in the error estimation for the objective in Eqn. (23).

6.2. Mesh adaptation in a multifidelity setting

With the requirements for high accuracy of the design, a fine discretization has to be chosen, so that the discrete optimization problem approximates the original continuous problem with required accuracy. Instead of optimizing on a mesh with fixed resolution, which would always require the highest fidelity for accurate calculations, the mesh is progressively refined as the optimization proceeds, resulting in a multifidelity optimization.

Rather than performing optimization and mesh adaptation sequentially, one after another, a concurrent framework is introduced. Two possible ways to incorporate the mesh adaptation and design optimization are considered here: optimization-driven adaptation and adaptation-driven optimization. In the former approach, the optimization tolerance at each fidelity is prescribed by the user. The objective function is first evaluated on a relatively coarse mesh, and then the error estimation and mesh adaptation are performed to control the discretization error to be below the optimization tolerance at the current fidelity. The allowable discretization error decreases as the optimization fidelity increases. The mesh adaptation techniques in this method use goal-oriented adaptation with Hessian-based anisotropy. For the latter approach, several mesh levels (degrees of freedom) are defined before the optimization. Again, we start with a fairly coarse mesh, and then the mesh is optimized for each design to achieve the best accuracy. Once the objective change or the gradient norm is smaller than the objective error estimate, the optimization terminates at the current cost level and the fidelity increases through mesh adaptation with a higher cost. This approach is designed for MOESS to optimize the mesh for a given cost at each fidelity. We refer to these methods as error-based or cost-based, depending on the information specified.

Compared with the fixed-fidelity optimization, unnecessarily fine meshes at the early stages of shape optimization are avoided in the two proposed multifidelity frameworks. Moreover, the areas that introduce most of the error may differ a lot for different designs during the optimization. Both approaches reduce the chance of over-refining areas that are not relatively important for the final design, which is important if the adaptation mechanics do not allow for coarsening. Compared with the multifidelity optimization without error estimation, the optimization tolerance and the error estimate are tightly coupled to actively control the optimization at each step and avoid the waste of low-fidelity convergence. Therefore, we expect that the two methods can effectively prevent over-optimizing on a coarse mesh, or over-refining on an unintended shape.

6.3. Consistent objective-sensitivity analysis

One should note that even at the same fidelity, the mesh is not necessarily fixed as in the traditional design method. Rather, the mesh is also adapted if needed, e.g. refined in the error-based method or optimized in the cost-based method to control the discretization error. Recall the discrete optimization problem in Eqn. (8), omitting the additional trim constraints for simplicity,

$$\begin{aligned} \min_{\mathbf{x}} \quad & J_h(\mathbf{U}_h(\mathbf{x}), \mathbf{x}) \\ \text{s.t.} \quad & \mathbf{R}_h(\mathbf{U}_h(\mathbf{x}), \mathbf{x}) = \mathbf{0} \end{aligned} \tag{40}$$

The above problem formulation infers a dependence on the discretization h , i.e. the computational mesh. One could prove that by refining the discretization, the discrete optimization problem converges to the continuous one. However, it should be mentioned that, the discrete optimization is an independent problem, which is characterized by the behavior of both the continuous problem and the discretization induced error.

$$J_h(\mathbf{u}_h, \mathbf{x}) = \mathcal{J}(\mathbf{u}, \mathbf{x}) + \mathcal{E}_h(\mathbf{u} - \mathbf{u}_h, \mathbf{x}) \quad \Rightarrow \quad \frac{dJ_h}{d\mathbf{x}} = \frac{d\mathcal{J}}{d\mathbf{x}} + \frac{d\mathcal{E}_h}{d\mathbf{x}} \quad (41)$$

In general, the discretized objective is incorrect compared to the exact continuous objective due to the numerical error induced by discretization. Consequently, the discretized objective gradient also involves the gradient of the error (assuming the error is a continuous function with respect to the design parameters), which is unique to each discretization. If the mesh is not adapted, then the numerical error may lead to a substantial deviation of the objective. Indeed, even if the mesh is adapted, but only on the objective and not its gradient, convergence to the true continuous objective may be hampered due to inaccuracy of the gradient approximation [22] – addressing such situations is beyond the scope of this paper. However, the discretization error does not affect the convergence of the discrete optimization problem, since the discretized gradient analysis is consistent with the discretized objective function. In fact, if exact differentiation is used in Eqn. (41), one obtains the exact gradient of the discretized objective functional.

In the proposed frameworks, inconsistent objective-sensitivity analysis may occur since the discretized optimization problem changes every time the mesh is adapted. Although all of the problems are approximations to the continuous optimization problem, each of them has its own behavior because of the embedded discretization error. There may not exist a feasible gradient-based update path between two design and discretization pairs, since the gradient depends on the discretization and hence cannot guide the update between different discretizations. Therefore, for the sake of consistent objective-gradient analysis, we update the design on the same mesh, and then perform mesh adaptation if needed, yielding a sequence of designs and discretizations, which converges to the optimal design of the continuous optimization problem,

$$\{J_h\} = \{J_{h_0}(\mathbf{x}_0), J_{h_0}(\mathbf{x}_1), J_{h_1}(\mathbf{x}_1), \dots, J_{h_{N-1}}(\mathbf{x}_N), J_{h_N}(\mathbf{x}_N)\} \quad \lim_{h_N \rightarrow 0} |J_{h_N}(\mathbf{x}_N) - \mathcal{J}(\mathbf{x}^*)| = 0 \quad (42)$$

6.4. Algorithm overview

The proposed optimization frameworks with error estimation and mesh adaptation are summarized in Algorithm 1 and Algorithm 2, using error-based and cost-based approaches respectively. Optimization tolerance levels or cost levels are specified by the user, driving the mesh adaptation to actively control the discretization errors. In this paper, we assume the error estimation is accurate enough to represent the “true” discretization error, which may be inappropriate when the adjoint is not well-resolved or when the problem is highly nonlinear. In practice, a safety factor η can be used to ensure the discretization error to be always below the optimization tolerance; $\eta = 1$ is adopted in this paper.

7. Results

In this section, we first demonstrate the importance of controlling the discretization error in optimization with a problem governed by one-dimensional advection diffusion equations. Then, two-dimensional airfoil shape optimization problems governed by the compressible Navier-Stokes equations are considered. The two methods proposed in Section 6 are applied in the airfoil optimization to compare their effectiveness in achieving optimality and their efficiency of reducing the error.

7.1. Scalar advection-diffusion

We now take a closer look at the example shown in Figure 1(b), the system is governed by the one-dimensional scalar advection-diffusion equation,

$$\begin{aligned} a \frac{\partial u}{\partial x} - \nu \frac{\partial^2 u}{\partial x^2} &= 0 & x \in [0, L] \\ u(0) &= 0 & u(L) = 1 \end{aligned} \quad (43)$$

Algorithm 1: optimization with error estimation and mesh adaptation (error-based)

input : initial design \mathbf{x}_0 , initial coarse mesh $\overline{\mathcal{T}}_h$, optimization tolerance levels $\tau_0, \tau_1, \dots, \tau_n$, safety factor $\eta \leq 1$
output: optimal design \mathbf{x}^* with controlled objective error $\mathcal{E}(J_h) \leq \tau_n$

```
1 for  $l = 0, 1, \dots, n$  do
2   set error tolerance as  $\mathcal{E}_l \leftarrow \eta\tau_l$ 
3   while not converged do ▷ optimization algorithm
4     compute objective function  $J_h(\mathbf{x}_l)$  and its error estimate  $\delta J_h(\mathbf{x}_l)$ 
5     while  $\delta J_h > \mathcal{E}_l$  do
6       adapt mesh  $\overline{\mathcal{T}}_h$ 
7     end
8     calculate objective gradient  $dJ_h/d\mathbf{x}$  and update design  $\mathbf{x}_l$  ▷ line search
9   end
10  finish optimization at level  $l$ ,  $\mathbf{x}_{l+1} = \mathbf{x}_l$ 
11 end
```

Algorithm 2: Optimization with error estimation and mesh adaptation (cost-based)

input : initial design \mathbf{x}_0 , initial coarse mesh $\overline{\mathcal{T}}_h$, cost levels C_0, C_1, \dots, C_n , safety factor $\eta \geq 1$
output: optimal design \mathbf{x}^* with optimal accuracy at given cost C_n

```
1 for  $l = 0, 1, \dots, n$  do
2   while not converged do ▷ optimization algorithm
3     optimize mesh  $\overline{\mathcal{T}}_h$  to minimize discretization error at fixed cost  $C_l$ 
4     calculate objective  $J_h(\mathbf{x}_l)$  and its error estimate  $\delta J_h(\mathbf{x}_l)$ 
5     set optimization tolerance  $\tau_l = \eta\delta J_h$ 
6     calculate objective gradient  $dJ_h/d\mathbf{x}$  and update design  $\mathbf{x}_l$  ▷ line search
7   end
8   finish optimization at level  $l$ ,  $\mathbf{x}_{l+1} = \mathbf{x}_l$ 
9 end
```

The optimization problem is formulated as seeking an optimal Peclet number defined as $Pe = aL/\nu$ to minimize the negative scalar gradient at a specified location,

$$\min_{Pe} -\frac{\partial u}{\partial x}\Big|_{x=0.76L} \quad (44)$$

The objective functional $\mathcal{J} = -\frac{\partial u}{\partial x}\Big|_{x=0.76L}$ is an implicit function of the design variable \mathbf{x} (Peclet number) defined by the underlying state equation. We discretize the continuous optimization problem using DG with approximation order $p = 2$. Three computational meshes with the same degrees of freedom are tested: a uniformly-distributed mesh, an isotropically-refined mesh, and an optimized mesh with MOESS, both starting from a coarser mesh. The discretized objective functionals on different meshes are shown in Figure 3(a). We see that the naive uniformly distributed mesh produces a spurious optimum besides the expected one due to the discretization error, while the adapted meshes are able to predict a reasonably-accurate objective functional over the entire design space. Therefore, the optimization performed on the uniform mesh can heavily depend on the starting point, especially for gradient-based methods. The optimizer may converge to the spurious local optimum if the descent direction is pointing to it. With the same degrees of freedom, the adapted meshes based on objective error are more robust to the starting point, since the objective shape over the design space is preserved by reducing the discretization error. However, the discretization error can hardly be eliminated: we can only converge to the optimum by some tolerance comparable to the objective error, which is always the case in most engineering applications. By zooming

into the region near the exact optimum, we can see that the uniform mesh has the highest error leading to a most inaccurate optimum. The optimized mesh tends to give better accuracy compared to the isotropically-adapted mesh. In order to more closely study the spurious optimum, we plot the state solution $u(x)$ over the computational domain on different meshes in Figure 3(b). We can observe a severe numerical oscillation near the location of interest for the solution on the uniform mesh, while the adapted meshes are made to predict an accurate objective, which reduces the possibility of the occurrence of spurious optima. The numerical oscillation shown in Figure 3(b) is only one of the possible sources that may cause spurious optima. Many physical features that are sensitive to discretization errors, such as boundary layers or shocks in flow problems, can also potentially create spurious optima.

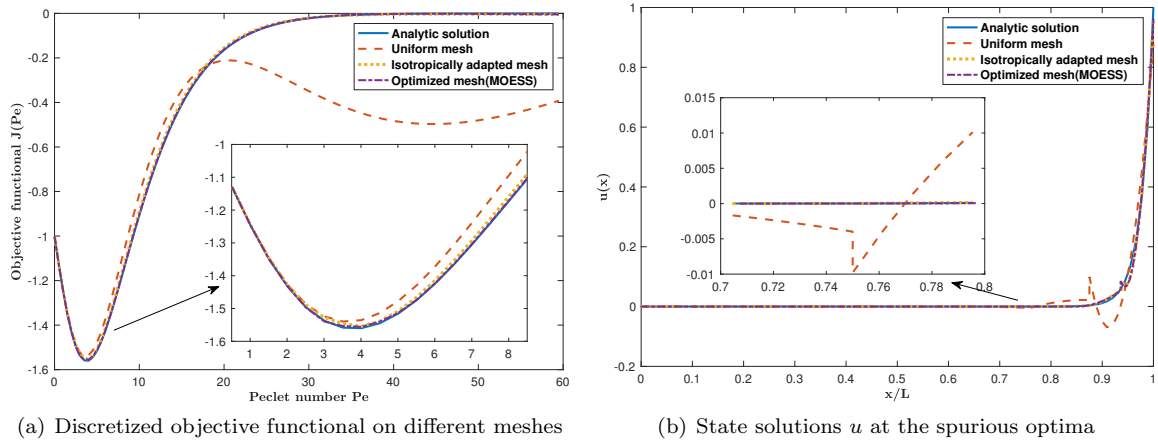


Figure 3: Optimization governed by 1D scalar advection-diffusion PDE on different meshes

A comparison of the optimization on these meshes starting from different Pelet numbers is listed in Table 1. We do not consider the error estimation and optimization coupling here since the error in the objective is fairly large on these coarse meshes, which can cause early convergence for the optimization. The optimization tolerance in these runs is set to be small enough compared to the discretization error. The optimization on the fixed uniform mesh converges to the spurious optimum when the initial point is close to it, whereas the adapted meshes produce more accurate results that are more robust to the starting point as expected. Furthermore, optimization with meshes produced by MOESS exhibits better accuracy since the optimized meshes tend to obtain a better approximation of the objective over the design space. Also, as shown in the table, if the optimizer converges to an inaccurate optimum that is close to the exact one, the error estimate gives a good prediction of the optimization accuracy. On the other hand, if the optimizer converges to a spurious optimum caused by the discretization error, the error estimates at the spurious optimum, while not very effective, are generally large enough to indicate the optimization is not converging correctly.

Table 1: Optimization results (advection-diffusion PDE) on different meshes (DG , $p = 2$, $N_e = 8$, $optimization\ tol = 10^{-10}$)

Initial design	Mesh	\mathbf{x}_h^*	$J_h(\mathbf{x}_h^*)$	$\delta J_h(\mathbf{x}_h^*)$	$\ \mathbf{x}_h - \mathbf{x}^*\ $	$\ J_h(\mathbf{x}_h^*) - \mathcal{J}(\mathbf{x}^*)\ $
$\mathbf{x}_0 = 40$	Uniform mesh	44.60287	-0.49700	-0.15555	40.79686	1.06446
	Iso-adapted mesh	3.72145	-1.54724	0.01309	0.08455	0.01423
	Optimized mesh	3.81001	-1.56218	-0.00063	0.00400	0.00072
$\mathbf{x}_0 = 20$	Uniform mesh	3.64579	-1.54037	0.02063	0.16022	0.02109
	Iso-adapted mesh	3.84404	-1.55622	0.00505	0.03803	0.00525
	Optimized mesh	3.81001	-1.56218	-0.00063	0.00400	0.00072

7.2. Airfoil shape optimization

The previous simple example using a one-dimensional advection-diffusion equation shows the importance of controlling the objective error in optimization governed by PDEs. However, it is more common in aerospace engineering applications to optimize the objective subject to some constraints, where the discretization error can affect both the objective and constraint outputs. The error in the constraints can often indirectly affect the calculation of the objective, hence the mesh should be adapted to predict both the objective and constraint outputs with appropriate accuracy. As a simple demonstration of the constrained optimization approach, we consider two-dimensional airfoil shape optimization problems in different flow regimes, which impose different governing equations to the optimization problem. The goal of the optimization is to seek an optimal airfoil shape and incidence angle to minimize the drag coefficient subject to fixed lift trim condition and the minimum volume constraint. We only consider the discretization errors in the drag and lift calculations, and the airfoil volume measurements are assumed to be exact. Furthermore, the lift trimming constraint tolerance is always set to be the same as the optimization tolerance.

The airfoil is parametrized with 10 Hicks-Henne basis functions [44], and cubic curved mesh elements are used to represent the boundary. At each line-search iteration in the optimization, the objective function needs to be re-evaluated, which requires a flow solution on the updated geometry, and hence a new mesh must be obtained every time. Regeneration of a mesh, especially for a complex geometry or with high resolution, could be time-consuming and non-trivial. Furthermore, mesh topology changes are not desirable for the purpose of objective-gradient consistency. Thus, an efficient way to update the computational mesh is needed. In this work, we use radial basis function (RBF) interpolation [45] to deform the mesh for small shape changes, and the linear elasticity equations [46] for large deformations. We first demonstrate our proposed error control approach with a simple smooth laminar flow problem. Then the two different adaptation strategies are compared in a more challenging inviscid transonic flow problem. Finally, a more practical turbulent transonic case is investigated.

7.2.1. Laminar, subsonic airfoil

A subsonic laminar flow case is first studied. The initial design is a NACA 0012 airfoil at zero angle of attack with a free-stream Mach number $M_\infty = 0.5$ and Reynolds number $Re = 5000$. Three optimization and error control strategies are investigated: fixed (highest) fidelity optimization without error estimation and mesh adaptation, multifidelity optimization with error estimation and mesh adaptation only on the objective (only the error in the drag coefficient, c_d , is controlled during the optimization), and the multifidelity optimization with meshes adapted on both the objective and constraints as proposed in Section 4 (total error in the objective is restricted). We only consider the effects of different error control strategies, thus only the goal-oriented mesh adaptation with Hessian-based anisotropy detection is used here. Different adaptation mechanics are compared in Section 7.2.2.

All of the optimization runs start from the same initial mesh consisting of 533 triangular elements with a DG $p = 1$ discretization. For the fixed-fidelity optimization, the mesh is first adapted to meet the objective tolerance, and then no more mesh adaptation occurs during the optimization. In contrast, in the other two methods, the mesh is adapted during the optimization, taking into account a changing error estimate. The initial symmetric airfoil should produce zero lift at zero incidence, and thus the initial condition is further from feasible if a higher target lift is specified, which means higher error of the outputs may appear during the optimization.

The target lift coefficient is set to be $c_l^* = 0.1$, with a minimum volume (area) constraint of 95% of the initial NACA 0012 airfoil volume, and the final optimization tolerance is 1×10^{-4} in the drag coefficient (i.e. one drag count). We should expect the final meshes to be comparable in size for all of the methods because of the same ultimate tolerance for the optimization. The meshes during the optimization, however, may be quite different. The fixed fidelity mesh and the meshes at the same intermediate optimization fidelity for the multifidelity methods are shown in Figure 4. This figure shows that multifidelity optimization significantly reduces the mesh size and computational resources during the early optimization iterations. The drag adjoint and lift adjoint (density component) are also shown with the meshes for the two multifidelity optimizations. Though the shape is not the same, the flow features and mesh sizes are similar. Most of the mesh adaptation

happens at the leading edge since the drag adjoint is the largest (and least resolved) near the airfoil nose, while the proposed adaptation based on combined error estimation, Figure 4(d), adapts more on the trailing edge, which is also important for accurate lift prediction.

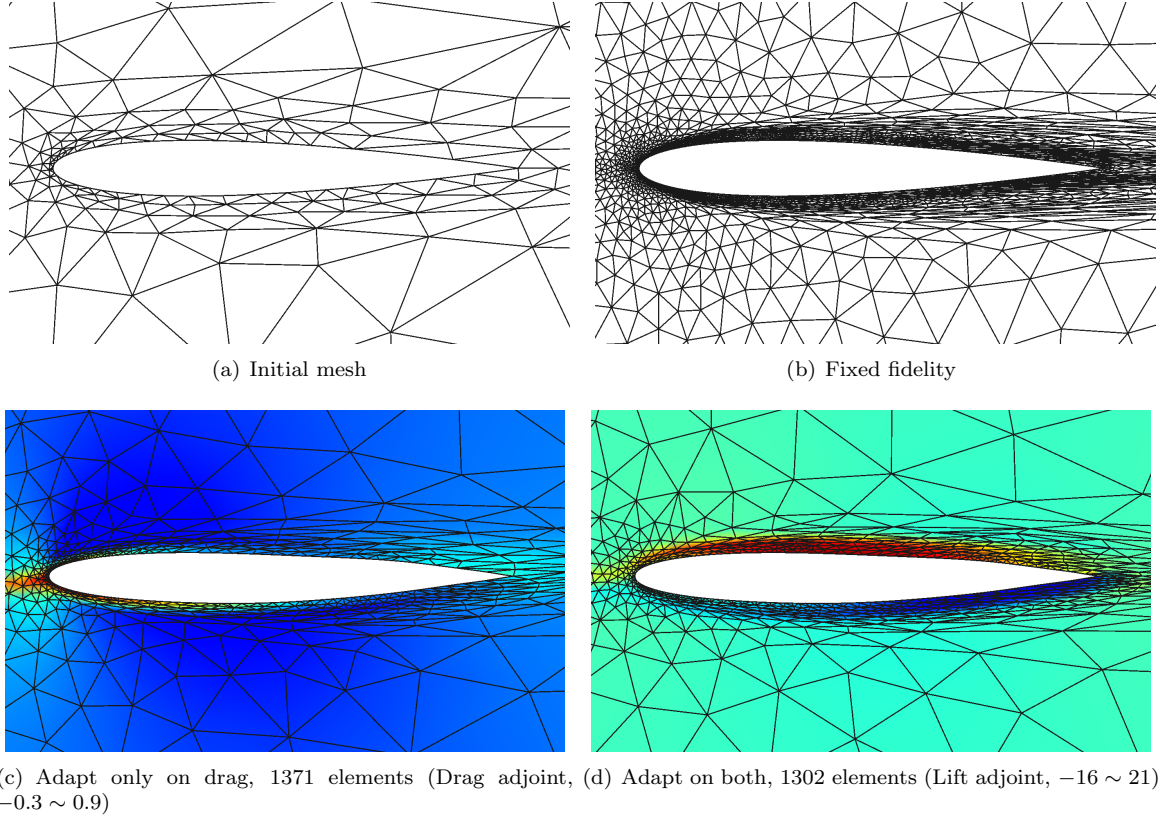
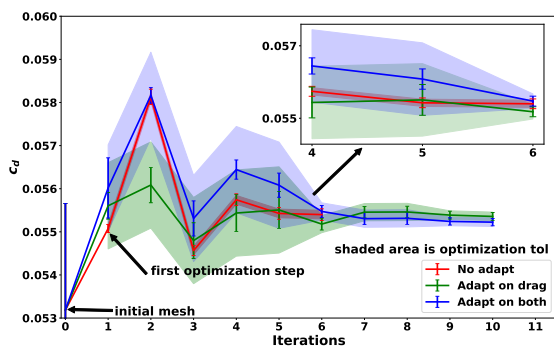


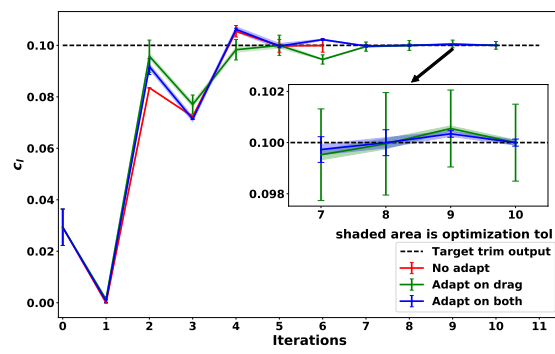
Figure 4: Initial mesh and intermediate meshes for different error control strategies.

The objective and constraint convergence histories are shown in Figure 5(a) and Figure 5(b). From the convergence plots, we see that the discretization error of the objective increases as the shape changes for the fixed-fidelity optimization without mesh adaptation. The objective error may be above the optimization tolerance during the optimization even though we start with a fairly fine mesh, and it converges to a noticeably different design compared to the results of the other two methods with error control, as shown in Figure 5(c). On the other hand, the objective error is always controlled to be below the optimization tolerance via mesh adaptation in the other two methods. However, the mesh adapted only on the objective has similar but slightly lower constraint error compared to the fixed mesh. The proposed method achieves the lowest constraint error by taking it into account in the error estimation and mesh adaptation, and of course requires higher computational cost at the highest fidelity, as shown in Figure 5(d). Furthermore, the fixed-fidelity optimization requires the most iterations on the highest fidelity (the number of objective evaluations would be even more because of the line search between each major iteration). Multifidelity frameworks benefit from fewer iterations on the finest mesh since they have better starting designs obtained from the coarser meshes.


Local Mach number contours of the final designs on the corresponding meshes achieved by all these three methods are given in Figure 6. Again, we can see that the mesh adapted on both the objective and constraint outputs has more refinement on the airfoil upper surface around the trailing edge, which is important for accurate lift calculation, while it has less of an effect on the drag prediction. We take the optimized shape as well as the final mesh produced by all these methods, then increase the approximation order to $p + 1$, and



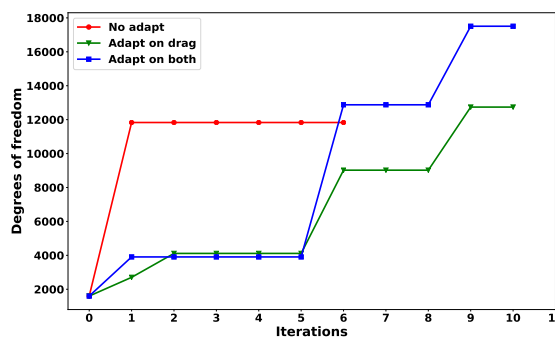
(a) Objective convergence history



(b) Constraint convergence history



(c) Initial and optimized shapes



(d) Mesh size evolution

Figure 5: Optimization history and final designs for different error controlling strategies (laminar, subsonic)

start a new optimization (fixed mesh) to get the “exact” optimization solutions on these meshes. Since the meshes are not infinitely fine, we do not expect the “exact” solutions to be the same. The main results for the optimization on different meshes are summarized in Table 2. Using Eqn. (23), we decompose the error of the optimal objective into two parts: pure objective error (third column) and the error due to inaccurate trim constraint (fourth column). If the mesh does not adapt to predict accurate constraints, the error in the constraints can indirectly affect the objective, which is comparable to the pure objective error as shown in the table. In the proposed method, both error sources are well controlled, as the total error estimate of the optimal objective is below the optimization tolerance. In this problem, we know that some refinements (e.g. the leading edge) for accurate drag prediction also improve the lift accuracy. However, the important areas for objective and constraints can be very different in some problems [33], and adapting only on the objective can lead to undesired designs due to inaccurate constraints. By comparing the difference between order p and $p + 1$ solutions, we also find the error estimate given by Eqn. (23) is pretty close to the “exact” error in this problem. However, the adjoint-based error estimates can be inaccurate on very coarse meshes or when the adjoints are not well resolved. In addition, the Lagrange multipliers extracted from the optimizer can be inaccurate when the objective and constraint gradients are not accurate or when the problem is highly nonlinear. We would not expect the error estimate in Eqn. (23) to be very accurate in these problems or situations, but it can still give the user an estimate of the error level which can be a guideline for the optimization (a safety factor η can be used to ensure the error to be sufficiently small compared to the optimization tolerance). The total optimization costs are also compared in Table 2 by scaling with the wall time used by the fixed fidelity optimization. Though the iterations on the finest meshes are reduced, we increase the total computational cost with error control and mesh adaptation due to additional adjoint and flow solves. Those extra costs required for more accurate design can possibly be reduced by more efficient adaptation techniques.

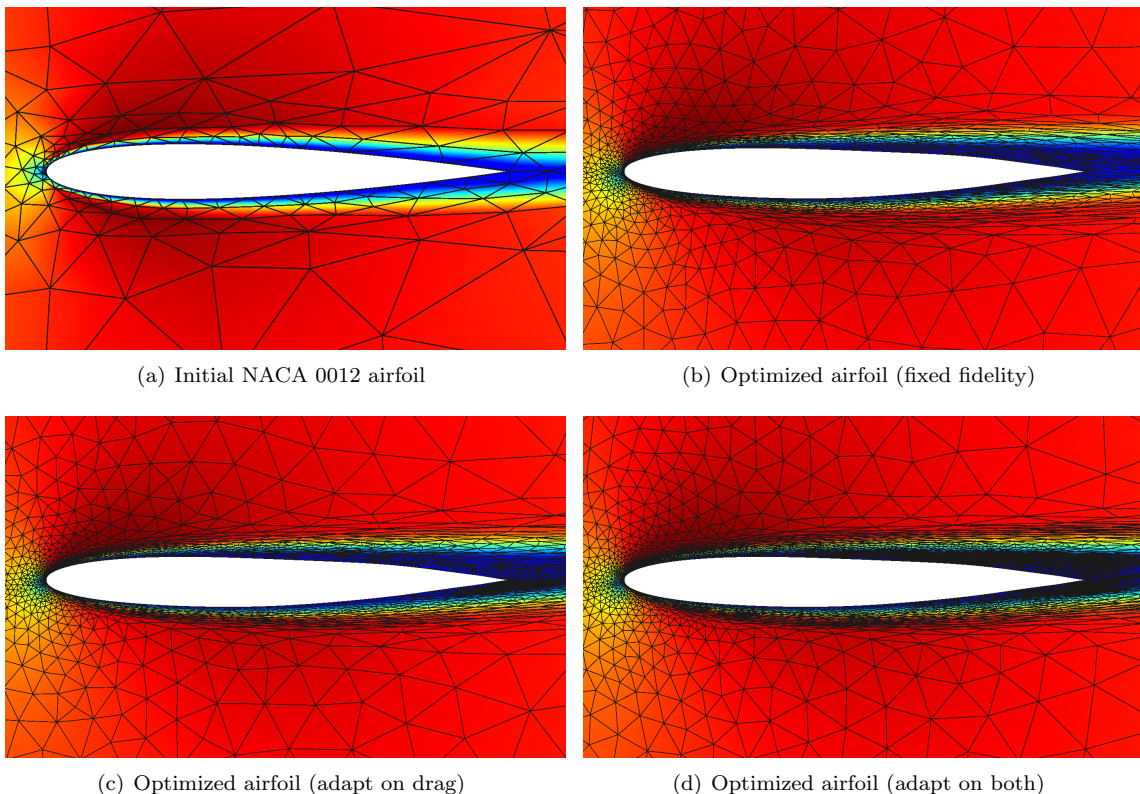


Figure 6: Local Mach number contour ($0 \sim 0.6$) for the initial and final designs (laminar, subsonic)

Table 2: Laminar airfoil optimization results on different meshes ($DG, p = 1$, $final\ optimization\ tol = 1e - 4$)

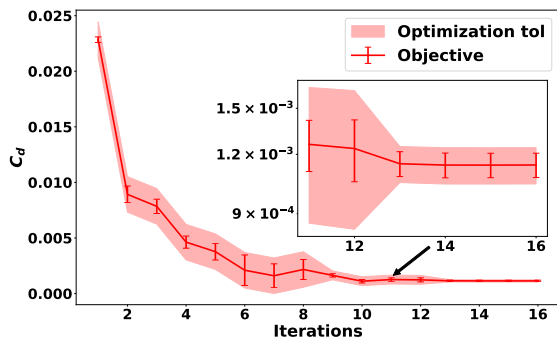
	$(J_{h,p}^{adapt})^*$	δJ^{adapt}	$\mu_h^T \delta \mathbf{J}^{trim}$	δJ_{opt}^{adapt}	$(J_{h,p+1}^{adapt})^*$	$\ (J_{h,p}^{adapt})^* - (J_{h,p+1}^{adapt})^* \ $	Cost
Fixed mesh	5.53993E-2	1.347E-4	1.456E-4	2.803E-4	5.51980E-2	2.013E-4	1.00
Adapt on drag	5.53572E-2	9.060E-5	8.490E-5	1.755E-4	5.51814E-2	1.758E-4	1.11
Adapt on both	5.52243E-2	8.402E-5	7.009E-6	9.103E-5	5.51327E-2	9.160E-5	2.56

7.2.2. Inviscid, transonic airfoil

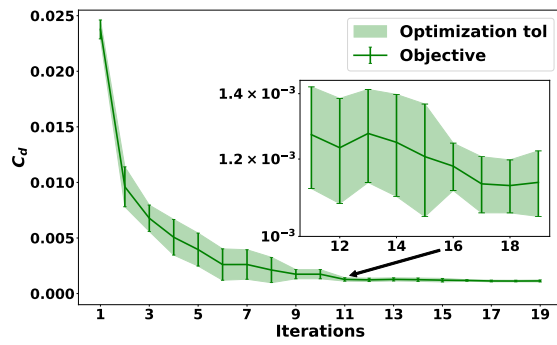
We next apply the new objective error estimation to an optimization problem based on an inviscid transonic flow at $M_\infty = 0.8$ around an initial NACA 0012 airfoil. The initial angle of attack is $\alpha = 1.25^\circ$, and the goal is to minimize the drag with a target lift coefficient of $c_l^* = 0.4$. Again the minimum volume of the airfoil is set as 95% of the initial design. Although the transonic flow around the original NACA 0012 airfoil features a strong shock on the upper surface near the trailing edge, we expect that the shape would be modified during the optimization such that the shock strength is weakened or the shock is completely removed. The flow features and the outputs of interest are highly related to the location and strength of the shock, which may change significantly during the optimization. Thus, error estimation and mesh adaptation become more crucial in this case. In this problem, the mesh adaptation is based on the objective error estimate including the constraint error effects. Meanwhile, shocks are one of the very common anisotropic features in flow problems, since the flow field changes a lot in the direction normal to the shock while remains similar along the shock direction. Therefore, two different anisotropic adaptation mechanics are tested here: goal-oriented mesh adaptation with Hessian-based anisotropy detection, and mesh optimization via error sampling and synthesis (MOESS). Due to the different adaptation mechanisms, the couplings between optimization and mesh adaptation are also different. The former uses an error-based multifidelity optimization framework, while the latter adopts a cost-based one, as described in Algorithm 1 and Algorithm 2, respectively.

We start with the same initial mesh used in the laminar case, the objective convergence histories for these two adaptation mechanics are shown in Figure 7(a) and Figure 7(b). The objective on the initial mesh is not shown in the plots for simplicity. Since the error estimates have already included the effect of the constraints, we do not show the constraints convergence plots here. The optimization with Hessian-based anisotropy detection is performed with user-specified multiple error levels with an ultimate tolerance of 1×10^{-4} for the objective. On the other hand, the optimization using MOESS starts at a fairly low cost level, and degrees of freedom are added once the optimization converges at current cost level. The optimization stops when the total objective error (setting to be the optimization tolerance) is below 1×10^{-4} . As shown in the convergence plots, these two methods have similar convergence for this optimization problem. Most of the drag reduction happens at the lowest fidelity, where the flow solve is very cheap, though quite inaccurate. Both methods converge fast at the highest fidelity by virtue of better starting shapes obtained from the lower fidelity. We obtain fairly close optimal objective values as shown in Table 3. Compared with the initial objective value $c_{d,0}$ (the initial drag coefficient in the table is obtained by a very fine mesh on the original NACA 0012 airfoil, not the first point in the convergence plot), the total drag is therefore reduced by around 95% for both methods. The optimized shapes on these two meshes are also very similar as shown in Figure 7(d). The final designs approach a flattened upper surface and a higher aft camber, resembling a super-critical airfoil. The mesh size evolution is plotted in Figure 7(c). The mesh size (measured by dof) of the MOESS approach is much smaller than the goal-oriented Hessian-based adaptation, which will be discussed in detail later.

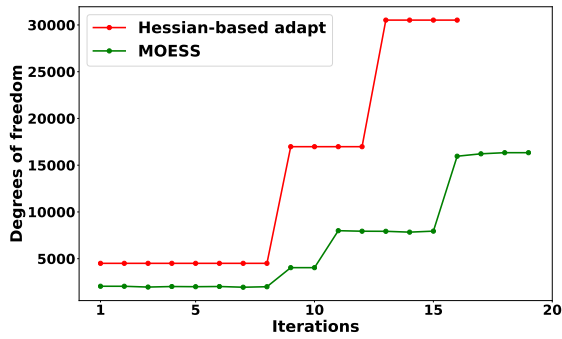
Mach number contours for the initial and final designs on the corresponding meshes are shown in Figure 8. The initial strong shock is significantly weakened on the optimized airfoils. We can still see weak discontinuities in the flow field, though, and the drag could be further reduced if more basis functions were used to parametrize the airfoil shape and even higher fidelity flow calculations were performed. One can also notice that for the final designs in Figure 8, the Hessian-based anisotropy detection method requires a mesh which is much finer compared to the mesh optimized by MOESS. In fact, the optimized mesh only



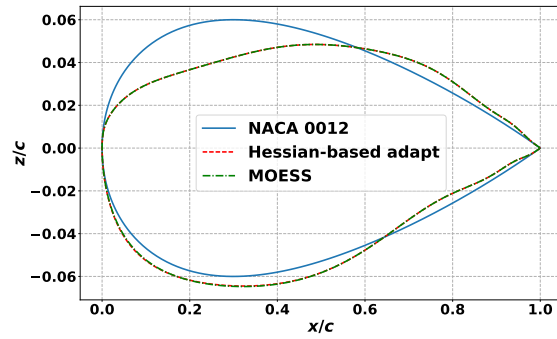
(a) Objective convergence history (Hessian-based adaptation)



(b) Objective convergence history (MOESS)



(c) Mesh size evolution



(d) Initial and optimized shapes

Figure 7: Optimization history and final designs for different adaptation mechanics (inviscid, transonic)

needs around half of the degrees of freedom required for the Hessian-based method, as shown in Figure 7(c). The total optimization cost with MOESS is thus also around half of the cost needed for the optimization with goal-oriented Hessian-based adaptation, as shown in Table 3. There are two main reasons making the optimization with MOESS more efficient. One is the adaptation mechanics, and the other is the optimization algorithm. In the Hessian-based anisotropy detection algorithm, areas where the scalar solution (Mach number) anisotropy is large may not always be important for accurate objective and constraint prediction. An output-based mesh anisotropy optimization algorithm, such as MOESS, can therefore better predict the most efficient anisotropy over the domain. Hence, we observe many refinements on the airfoil surface in Figure 8(b) which cannot be seen in Figure 8(c). With regard to the optimization algorithm, error-based multifidelity optimization can be inefficient, since the mesh gets refined every time the error is higher than the optimization tolerance. As the shape changes, the mesh may get adapted in many areas that are important for different shapes, accumulating over-refinement during the optimization. On the other hand, the mesh is not required to be adapted as long as the error estimates are still below the tolerance, even when the shape and flow field change. This effect can be found in the optimization in this example. Some intermediate optimization steps of these two different methods are shown in Figure 9. In the 9th optimization step with goal-oriented Hessian-based adaptation, we can see that there is refinement aligned with the weak shock near the trailing edge in Figure 9(a). This refinement persists in the following optimization steps even though the weak shock has moved forward, since the error estimates are still below the optimization tolerance. Thus we lose some of the approximation capacity of these degrees of freedom. On the contrary, for the shape optimization combined with mesh optimization, we have refinements that always align to the weak shock on the top surface, and we can also observe that more anisotropic elements are added to the airfoil surface to improve the output accuracy, with the total degrees of freedom remaining the same. Therefore, for optimization with dramatic design changes or with complex systems involving strong anisotropy, MOESS will be more efficient and effective.

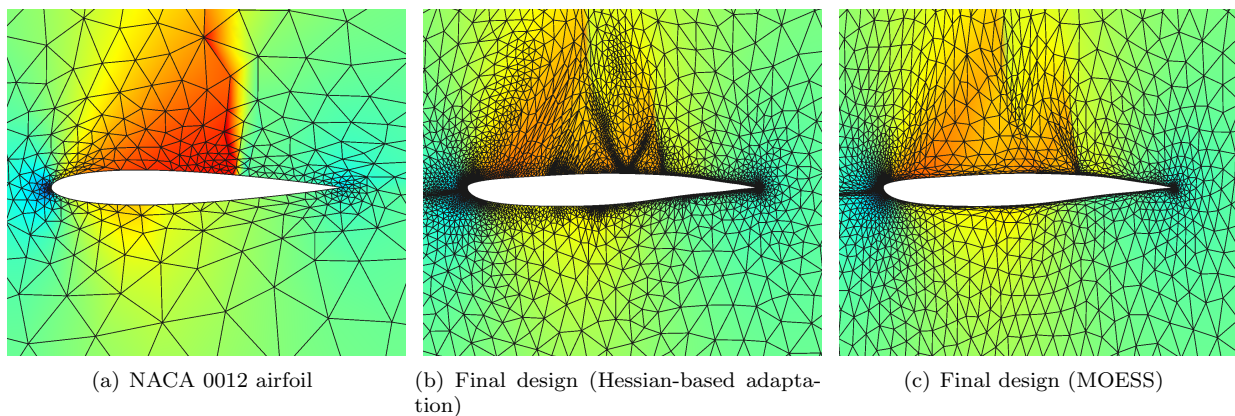




Figure 8: Local Mach number ($0 \sim 1.6$) for the initial and final designs (inviscid, transonic)

Table 3: Inviscid transonic airfoil optimization results on different meshes ($DG, p = 1$)

	$c_{d,0}$	$c_{d,opt}$	$\delta c_{d,opt}$	$\ c_{d,opt} - c_{d,0}\ /c_{d,0}$	Cost
Hessian-based	2.242E - 2	1.140E-3	6.708E-5	94.92% \pm 0.30%	1.00
MOESS		1.136E-3	8.752E-5	94.93% \pm 0.39%	0.53



(a) Goal-oriented Hessian-based adaptation, 9th, 10th and 12th optimization step from left to right



(b) MOESS, 16th, 17th and 19th optimization step from left to right

Figure 9: Intermediate meshes for different adaptation mechanics (inviscid, transonic)

7.2.3. Turbulent, transonic high-lift airfoil

The final problem considered in this paper is a more sophisticated turbulent case: optimization of a transonic airfoil with a high lift demand, starting with the RAE 2822 airfoil. The initial angle of attack is 2.79° with a freestream Mach number of $M_\infty = 0.734$ and Reynolds number $Re = 6.5 \times 10^6$. The target lift coefficient is set to be $c_l^* = 0.824$, and the airfoil volume is restricted to be no smaller than the initial value. The starting mesh for the turbulent case consists of 1448 elements, which is finer than the initial mesh used for the laminar and transonic runs. The flow simulation is fully turbulent with an approximation order of $p = 2$ in this case.

The initial mesh and the meshes used in the optimization are summarized in Figure 10. For the turbulent case, most of the adaptation focuses on resolving the boundary layer, which is highly anisotropic. Therefore, many degrees of freedom are put into the boundary layer with anisotropic elements around the airfoil boundary. Once the boundary layer is properly resolved, the objective error drops down very quickly. Then, more degrees of freedom are added to the upper surface to resolve the flow field changes along the flow direction.

The mesh size and objective are collected at each optimization step as shown in Figure 11. Here we assume that at a given fixed cost, the objective error on optimized meshes should be close if the physics are similar. This can also be observed in Figure 11(a). However we do see some oscillations of the error at coarse and medium meshes, which are caused by the shape changes and hence the different physics (shock strength and location). At the highest fidelity we expect only small changes in the shape, and the objective error remains almost the same, which means that both the shape and the mesh converge to an optimum. The initial and optimized airfoils are compared in Figure 12. As shown in Figure 12(c), the optimization flattens the upper surface near the forward section, while curving the lower surface and increasing the thickness in the aft section. The curvature reduction on the top surface is trying to smooth the flow acceleration region to weaken the shock. The thickened lower surface and aft section are required to maintain the lift and area constraints. This can be further observed in the Mach contours shown in Figure 12(a) and Figure 12(b), where the initial strong shock is significantly reduced. Hence the strong discontinuity is absent in the airfoil surface pressure distribution in Figure 12(d). Therefore the final design yields a drag coefficient of $c_{d,\text{opt}} = 105.06 \pm 0.28$ counts achieving a $43.73\% \pm 0.15\%$ drag reduction compared with the initial design, which produces a drag of $c_{d,0} = 186.71$ counts (the initial drag coefficient is obtained with an optimized finer mesh on the original design). This drag reduction is primarily due to an optimized pressure distribution, as the skin friction, shown in Figure 12(e), increases slightly on the upper surface, in the region where the shock is eliminated.

8. Conclusion

In most aerodynamic optimizations, we work with discretized governing flow equations. Thus, discretization-induced numerical error should be carefully controlled to ensure convergence to the “true” optimal design at a prescribed fidelity. Without properly controlling this error, the optimizer may arrive at a sub-optimal design or even at an incorrect spurious optimum with inaccurate information provided by the flow solver and gradient analysis, as shown in the test cases.

In this work, we present frameworks that integrate output-based error estimation and mesh adaptation with a traditional gradient-based algorithm. A coupled adjoint is also introduced, offering a way to include the constraints error into the objective error estimation. The multifidelity optimization approach consists of progressive refinement of the computational mesh and is capable of preventing over-optimizing and over-refining. The mesh adaptation (fidelity increase) is tightly coupled with the optimization algorithm either with an error-based or a cost-based strategy. Design optimization with mesh optimization via error sampling and synthesis (MOESS) is shown to be more efficient and effective by yielding meshes that optimize the approximation potential at a given computational cost. This benefit can become more significant when higher fidelity is required, or when more highly anisotropic physics govern the system.

With more judicious considerations of the objective functions and constraints, and additional design parameters, the new methods can provide realistic configurations in practical design scenarios. The fidelity

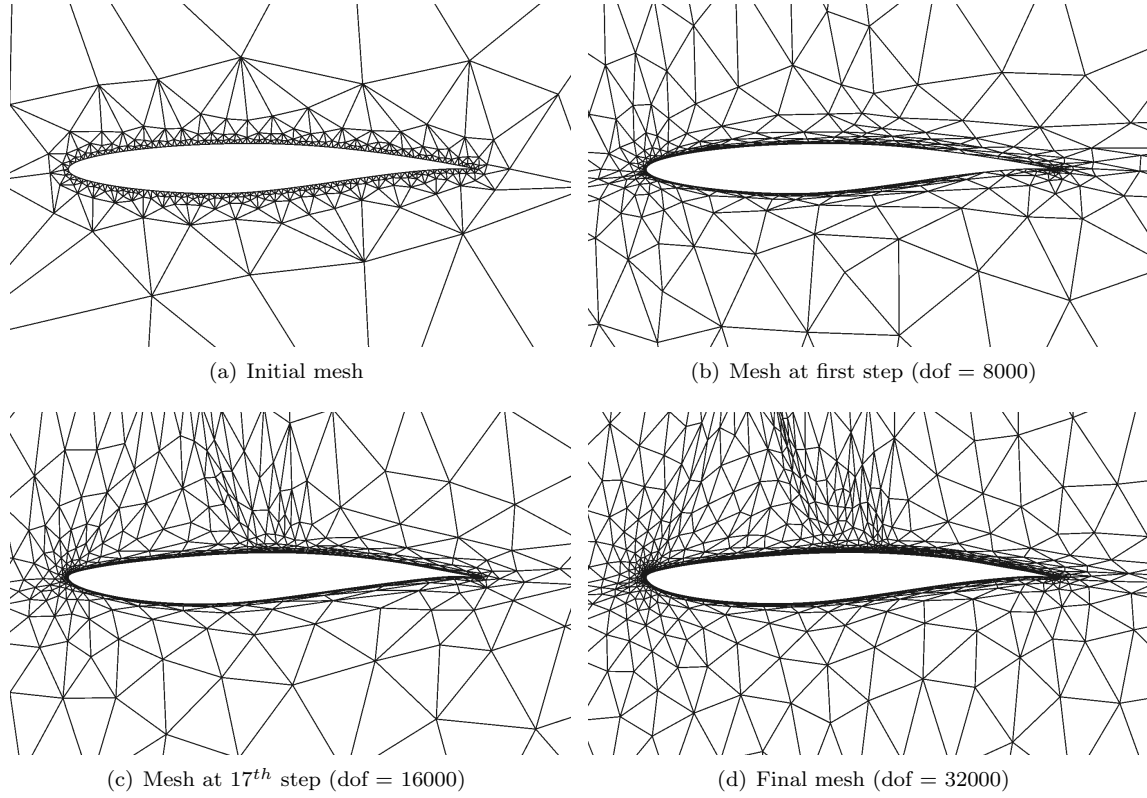


Figure 10: Initial mesh and intermediate meshes during the optimization (turbulent, transonic)



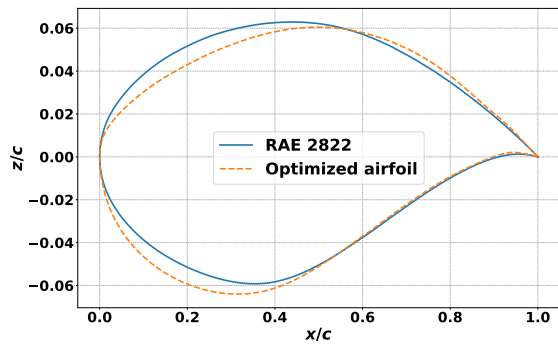


Figure 11: Objective convergence history and mesh evolution (turbulent, transonic)




(a) RAE 2822 airfoil


(b) Optimized airfoil



(c) Initial and optimized shapes



(d) Pressure distribution



(e) Skin friction coefficient distribution

Figure 12: Local Mach number ($0 \sim 1.3$), pressure, and skin friction distributions for the initial and final designs (turbulent, transonic)

increase is presently driven by an adaptation tolerance or a computational cost that increases or decreases by a fixed factor each time, or has to be specified by the user. However, for more practical problems, without a priori knowledge of the objective convergence, an improved and automated fidelity increase strategy should be developed to fulfill the potential of the present optimization framework to increase the accuracy and efficiency for aerodynamic optimization problems. Furthermore, only mesh adaptation (h -adaptation) is considered here to control the discretization error. More efficient adaptation mechanics such as approximation order increment (p -adaptation), and combinations (hp -adaptation) can also be applied to the proposed methods in the future.

Acknowledgments

The authors acknowledge the support of the Boeing Company, with technical monitor Dr. Mori Mani, and the Department of Energy under grant DE-FG02-13ER26146/DE-SC0010341.

References

- [1] J. R. Martins, P. Sturdza, J. J. Alonso, The complex-step derivative approximation, *ACM Transactions on Mathematical Software (TOMS)* 29 (3) (2003) 245–262. doi:10.1145/838250.838251.
- [2] A. Griewank, A. Walther, *Evaluating Derivatives: Principles and techniques of algorithmic differentiation*, Society for Industrial and Applied Mathematics, 2008. doi:10.1137/1.9780898717761.
- [3] J. R. Martins, J. J. Alonso, J. J. Reuther, A coupled-adjoint sensitivity analysis method for high-fidelity aero-structural design, *Optimization and Engineering* 6 (1) (2005) 33–62. doi:10.1023/B:OPTE.0000048536.47956.62.
- [4] A. Jameson, Aerodynamic design via control theory, *Journal of Scientific Computing* 3 (3) (1988) 233–260. doi:10.1007/BF01061285.
- [5] J. J. Reuther, A. Jameson, J. J. Alonso, M. J. Rimlinger, D. Saunders, Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers, part i, *Journal of Aircraft* 36 (1) (1999) 51–60. doi:10.2514/2.2413.
- [6] W. K. Anderson, V. Venkatakrishnan, Aerodynamic design optimization on unstructured grids with a continuous adjoint formulation, *Computers & Fluids* 28 (4-5) (1999) 443–480. doi:10.1016/S0045-7930(98)00041-3.
- [7] M. B. Giles, N. A. Pierce, An introduction to the adjoint approach to design, *Flow, Turbulence and Combustion* 65 (3-4) (2000) 393–415. doi:10.1023/A:1011430410075.
- [8] O. C. Zienkiewicz, J. Z. Zhu, A simple error estimator and adaptive procedure for practical engineering analysis, *International Journal for Numerical Methods in Engineering* 24 (2) (1987) 337–357. doi:10.1002/nme.1620240206.
- [9] I. Babuška, A. Miller, A feedback finite element method with a posteriori error estimation, part i: The finite element method and some basic properties of the a posteriori error estimator, *Computer Methods in Applied Mechanics and Engineering* 61 (1) (1987) 1–40. doi:10.1016/0045-7825(87)90114-9.
- [10] O. Zienkiewicz, J. Zhu, Adaptivity and mesh generation, *International Journal for Numerical Methods in Engineering* 32 (4) (1991) 783–810. doi:10.1002/nme.1620320409.
- [11] R. Mueller, D. Gross, G. Maugin, Use of material forces in adaptive finite element methods, *Computational Mechanics* 33 (6) (2004) 421–434. doi:10.1007/s00466-003-0543-z.
- [12] N. Kikuchi, K. Y. Chung, T. Torigaki, J. E. Taylor, Adaptive finite element methods for shape optimization of linearly elastic structures, in: *The Optimum Shape*, Springer, 1986, pp. 139–169. doi:10.1007/978-1-4615-9483-3_6.
- [13] N. Banichuk, F. Barthold, A. Falk, E. Stein, Mesh refinement for shape optimization, *Structural Optimization* 9 (1) (1995) 46–51. doi:10.1007/BF01742644.
- [14] A. Schleupen, K. Maute, E. Ramm, Adaptive fe-procedures in shape optimization, *Structural and Multidisciplinary Optimization* 19 (4) (2000) 282–302. doi:10.1007/s001580050125.
- [15] R. Becker, R. Rannacher, An optimal control approach to a posteriori error estimation in finite element methods, *Acta Numerica* 10 (2001) 1–102. doi:10.1017/S0962492901000010.
- [16] R. Hartmann, P. Houston, Adaptive discontinuous galerkin finite element methods for the compressible euler equations, *Journal of Computational Physics* 183 (2) (2002) 508–532. doi:10.1006/jcph.2002.7206.
- [17] D. A. Venditti, D. L. Darmofal, Grid adaptation for functional outputs: application to two-dimensional inviscid flows, *Journal of Computational Physics* 176 (1) (2002) 40–69. doi:10.1006/jcph.2001.6967.
- [18] K. J. Fidkowski, D. L. Darmofal, Review of output-based error estimation and mesh adaptation in computational fluid dynamics, *AIAA Journal* 49 (4) (2011) 673–694. doi:10.2514/1.J050073.
- [19] J. Lu, An a posteriori error control framework for adaptive precision optimization using discontinuous galerkin finite element method, Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA, available: <http://hdl.handle.net/1721.1/34134>, (2005).
- [20] M. Nemeč, M. Aftosmis, Output error estimates and mesh refinement in aerodynamic shape optimization, in: *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2013, p. 0865. doi:10.2514/6.2013-865.

- [21] D. Li, R. Hartmann, Adjoint-based airfoil optimization with discretization, Adjoint error control, *International Journal for Numerical Methods in Fluids* 77 (1) (2015) 1–17. doi:10.1002/flid.3971.
- [22] J. E. Hicken, J. J. Alonso, Pde-constrained optimization with error estimation and control, *Journal of Computational Physics* 263 (2014) 136–150. doi:10.1016/j.jcp.2013.12.050.
- [23] G. Chen, K. Fidkowski, Airfoil shape optimization using output-based adapted meshes, in: 23rd AIAA Computational Fluid Dynamics Conference, 2017, p. 3102. doi:10.2514/6.2017-3102.
- [24] M. Yano, D. L. Darmofal, An optimization-based framework for anisotropic simplex mesh adaptation, *Journal of Computational Physics* 231 (22) (2012) 7626–7649. doi:10.1016/j.jcp.2012.06.040.
- [25] S. Boyd, L. Vandenberghe, *Convex optimization*, Cambridge University Press, 2004. doi:10.1017/CB09780511804441.
- [26] G. Biros, O. Ghattas, Parallel lagrange-newton-krylov-schur methods for pde-constrained optimization, part i: The krylov-schur solver, *SIAM Journal on Scientific Computing* 27 (2) (2005) 687–713. doi:10.1137/S106482750241565X.
- [27] S. R. Allmaras, F. T. Johnson, Modifications and clarifications for the implementation of the spalart-allmaras turbulence model, in: Seventh international conference on computational fluid dynamics (ICCFD7), 2012, pp. 1–11.
- [28] C. E. Baumann, J. T. Oden, A discontinuous hp finite element method for convection-diffusion problems, *Computer Methods in Applied Mechanics and Engineering* 175 (3) (1999) 311 – 341. doi:10.1016/S0045-7825(98)00359-4.
- [29] P. Houston, E. Süli, hp-adaptive discontinuous galerkin finite element methods for first-order hyperbolic problems, *SIAM Journal on Scientific Computing* 23 (4) (2001) 1226–1252. doi:10.1137/S1064827500378799.
- [30] L. Wang, D. J. Mavriplis, Adjoint-based h-p adaptive discontinuous galerkin methods for the 2d compressible euler equations, *Journal of Computational Physics* 228 (20) (2009) 7643 – 7661. doi:10.1016/j.jcp.2009.07.012.
- [31] P. L. Roe, Approximate riemann solvers, parameter vectors, and difference schemes, *Journal of Computational Physics* 43 (2) (1981) 357–372. doi:10.1016/0021-9991(81)90128-5.
- [32] F. Bassi, S. Rebay, Gmres discontinuous galerkin solution of the compressible navier-stokes equations, in: *Discontinuous Galerkin Methods*, Springer, 2000, pp. 197–208. doi:10.1007/978-3-642-59721-3_14.
- [33] B. A. Rothacker, M. Ceze, K. Fidkowski, Adjoint-based error estimation and mesh adaptation for problems with output constraints, in: 32nd AIAA Applied Aerodynamics Conference, 2014, p. 2576. doi:10.2514/6.2014-2576.
- [34] A. Loseille, F. Alauzet, Continuous mesh framework part i: Well-posed continuous interpolation error, *SIAM Journal on Numerical Analysis* 49 (1) (2011) 38–60. doi:10.1137/090754078.
- [35] F. Hecht, Bamg: Bidimensional anisotropic mesh generator, User guide, INRIA, Rocquencourt (1998).
- [36] X. Pennec, P. Fillard, N. Ayache, A riemannian framework for tensor computing, *International Journal of Computer Vision* 66 (1) (2006) 41–66. doi:10.1007/s11263-005-3222-z.
- [37] M. Castro-Díaz, F. Hecht, B. Mohammadi, O. Pironneau, Anisotropic unstructured mesh adaption for flow simulations, *International Journal for Numerical Methods in Fluids* 25 (4) (1997) 475–491.
- [38] D. A. Venditti, D. L. Darmofal, Anisotropic grid adaptation for functional outputs: application to two-dimensional viscous flows, *Journal of Computational Physics* 187 (1) (2003) 22–46. doi:10.1016/S0021-9991(03)00074-3.
- [39] A. Loseille, F. Alauzet, Continuous mesh framework part ii: Validations and applications, *SIAM Journal on Numerical Analysis* 49 (1) (2011) 61–86. doi:10.1137/10078654X.
- [40] K. J. Fidkowski, D. L. Darmofal, A triangular cut-cell adaptive method for high-order discretizations of the compressible navier-stokes equations, *Journal of Computational Physics* 225 (2) (2007) 1653–1672. doi:10.1016/j.jcp.2007.02.007.
- [41] K. Fidkowski, A local sampling approach to anisotropic metric-based mesh optimization, in: 54th AIAA Aerospace Sciences Meeting, 2016, p. 0835. doi:10.2514/6.2016-0835.
- [42] D. Kraft, A software package for sequential quadratic programming, Tech. Rep. DFVLR-FB 88-28, DLR German Aerospace Center–Institute for Flight Mechanics, Köln, Germany (1988).
- [43] C. G. Broyden, A class of methods for solving nonlinear simultaneous equations, *Mathematics of Computation* 19 (92) (1965) 577–593. doi:10.1090/S0025-5718-1965-0198670-6.
- [44] R. M. Hicks, P. A. Henne, Wing design by numerical optimization, *Journal of Aircraft* 15 (7) (1978) 407–412. doi:10.2514/3.58379.
- [45] S. Jakobsson, O. Amoignon, Mesh deformation using radial basis functions for gradient-based aerodynamic shape optimization, *Computers & Fluids* 36 (6) (2007) 1119 – 1136. doi:10.1016/j.compfluid.2006.11.002.
- [46] J. E. Hicken, D. W. Zingg, Aerodynamic optimization algorithm with integrated geometry parameterization and mesh movement, *AIAA Journal* 48 (2) (2010) 400–413. doi:10.2514/1.44033.