



DATA-DRIVEN AERODYNAMIC SHAPE OPTIMIZATION AND MULTI-FIDELITY DESIGN EXPLORATION USING CONDITIONAL DIFFUSION-BASED GEOMETRY SAMPLING METHOD

Aobo Yang¹, Jinouwen Zhang², Jichao Li³ & Rhea Liem¹

¹The Hong Kong University of Science and Technology, Hong Kong SAR, China

²Shanghai Artificial Intelligence Laboratory, China

³Institute of High Performance Computing, A*STAR, Singapore

Abstract

Data-driven approaches have been shown to accelerate the process of high-fidelity design optimization, which has been successfully demonstrated in aerodynamic shape optimization (ASO) to support efficient and high-performance aircraft design. However, the high computational cost associated with generating aerodynamic data to build the required surrogate model contradicts our initial goal of achieving high-fidelity design with minimal computational expense. In this study, we exploit the capabilities of diffusion models in capturing complex high-dimensional data distributions to present an innovative conditional diffusion-based geometry sampler. Our research has led to the development of a comprehensive framework for high-fidelity data-driven ASO with reduced computational costs. This framework leverages multi-layer perception (MLP) as surrogate model and diffusion model as geometry sampler to achieve diverse and realistic geometric shapes while providing fine-grained control over specific properties. Moreover, we employ a conditional diffusion-based geometry sampling method to significantly reduce the training set requirement for the surrogate model construction, achieving a 50.5% decrease of the number of required data points. The effectiveness and capabilities of our proposed geometry sampling are validated using a high-dimensional dataset, demonstrating its potential to achieve high-fidelity ASO and multi-fidelity design exploration in a computationally efficient manner. In particular, the proposed geometry sampling method addresses several key aspects to support high-dimensional data-driven ASO and multi-fidelity design explorations, which include modeling the high-dimensional data distributions, sampling with time step and providing value function as the sampling guidance. This approach presents a comprehensive solution for effectively addressing the challenges associated with generating aerodynamic performance in high-dimensional and high-fidelity data. The associated experiments show this approach has the potential to serve the rapid high-fidelity design variations in modern aviation industry.

Keywords: Aerodynamic shape optimization, conditional diffusion model, data-driven and multi-fidelity design

1 Introduction

In recent years, data-driven design has become an effective approach in modern aircraft design, thanks to the availability of abundant data. Fueled by increasing volumes of data, the emerging machine learning methods can serve as data-driven optimization techniques capable of handling high-dimensional, non-convex, constrained, and multi-objective optimization problems [1]. Data-driven

approaches have shown that data, along with the more rigorous physics-based models, help make fast and reliable aerodynamic shape optimization (ASO) design possible [2, 3]. In the context of ASO, the data typically come from aerodynamic analysis results obtained by performing computational fluid dynamics (CFD) simulations. It is worth noting that the accuracy of models and optimization results depends on the level of fidelity of data. In particular, high-fidelity data (e.g., aerodynamic analysis data from using fine mesh) are required to ensure a higher accuracy, at the expense of high computational cost of data acquisition. Indeed, the quality of the data directly determines the accuracy of surrogate model predictions and the computational costs associated with generating these aerodynamic data for high-fidelity simulations. In this study, we present a new approach to sample the geometry shapes, which is called the conditional diffusion-based geometry sampling. This method enables generating data that are well-aligned with the desired design conditions, to yield a well-defined design space. By leveraging the benefits of the diffusion model, which is a time-dependent stochastic process capable of capturing complex and multimodal distributions [4], we effectively enhance the quality of our data. This, in turn, improves the accuracy of surrogate model predictions while substantially reducing the computational costs associated with generating such data. Building upon this premise, we proceed to review the related works in two subsequent sections. In Section 1.1, we review the data-driven design techniques and multi-fidelity design techniques that use surrogate model. In Section 1.2, we review the related sampling methods that define the data space of high-fidelity design. Specifically, we introduce the background of the conditional diffusion model and discuss why it is a suitable sampling method for such complex design challenges.

1.1 Data-driven and multi-fidelity design

In this paper, we mainly focus on data-driven techniques that use surrogate models, instead of relying only on full simulations, for design optimization. Compared to the conventional CFD-based optimization, the data-driven approach relies on a training set to capture various design variations and utilizes surrogate models to predict possible designs [5]. This methodology is commonly employed in optimization problems that involve substantial geometric variations, such as conceptual design, and is often regarded as a way to reduce dependence on computationally expensive CFD simulations. Such optimizations can be performed by either gradient-free or gradient-based algorithms. Gradient-free algorithms are useful when gradients are not available. One of their advantages is that they do not assume function continuity [6]. In large-scale CFD-based optimization, it has been noted that for smaller-sized problems, typically with 30 or fewer variables, gradient-free methods can be effective in finding a solution [7]. Similarly, combination of surrogate model for coefficients estimation and gradient-free algorithms have shown effectiveness in a series of studies [8, 9, 10, 11]. However, when it comes to high-dimensional problems, gradient-free methods often become less effective due to the increased computational costs associated with exploring a larger design space. Thus, gradient-based algorithm becomes an effective option for large-scale design variables.

Surrogate-based ASO has been commonly used to achieve high-fidelity design optimization of wings. The main advantage of using a surrogate-based ASO is the notable reduction in computational cost by replacing the computationally expensive high-fidelity CFD simulations with fast-to-evaluate surrogate models. Previous research in surrogate-based ASO usually uses gradient-free optimizers (i.e. genetic algorithms [GA] [12, 13, 14], Bayesian optimization [BO] [15, 16, 17], etc.) for two main reasons. First, the surrogate models (such as kriging and radial basis functions) do not always provide accurate, reliable gradient information. Second, gradient-free methods can effectively navigate the optimization landscape without the need for gradient computations. However, gradient-free optimizers generally struggle with high-dimensional data because the size of the search space grows exponentially with the number of dimensions, making it increasingly difficult for these algorithms to effectively explore the design space and converge to the global optimum in a reasonable amount of time. Thus, the surrogate-based ASO with gradient-free optimizers can typically only handle low-dimensional data. The insufficient data dimensionality inevitably leads to a lack of representational capacity for the characteristics of three-dimensional wing geometries.

In the realm of multi-fidelity design, a series of studies have utilized surrogate models to explore the relationship between different levels of data fidelity. This approach aims to achieve high-fidelity design optimization by leveraging a large quantity of low-fidelity data in conjunction with a smaller set of high-fidelity data. Among the various multi-fidelity techniques that use surrogate model, one of the most typical approaches for multi-fidelity modeling is Gaussian process regression (GPR) [18], which is also known as co-kriging [19]. GPR has the advantage of modeling data distributions and providing interpretability on regression. However, the conventional GPR calculation involves an inversion of a covariance matrix, which has a computational complexity of $O(N^3)$ for N data points. This can be computationally expensive and impractical for large datasets. Furthermore, the covariance matrix of GPR becomes sparsely populated, making it difficult to estimate the underlying function accurately. Thus, GPR suffers from the *curse of dimensionality*, where the performance degrades as the input dimensionality increases. With this consideration, deep neural networks (DNN) architecture is deemed more appropriate as it is highly scalable and can handle large-scale datasets with high-dimensional features. The flexible architecture makes DNN capable of capturing and modeling highly non-linear relationships in data [20]. This partially solves the issue of traditional Gaussian processes being unable to handle large amounts of data and high-dimensional features. As for using NN as a surrogate model, Meng and Karniadakis [21] proposed a composite neural networks that can be trained using multi-fidelity data. They used three NNs to capture correlations between low-fidelity and high-fidelity data that yielded accurate results; the method was then extended to physics-informed neural networks (PINNs) for various applications. Guo *et al.* [22] integrated the strengths of the approaches described above and proposed the multistep NNs for multifidelity regression. The multifidelity NN produces outputs with comparable accuracy to those of the expensive full-order model, using only a small number of full-order evaluations and a larger number of inexpensive evaluations from a reduced order model that are less accurate. There have also been works on utilizing multifidelity NN in ASO, such as in the designs of DLR-F4 wing-body configuration [23]. In our previous investigation on data-driven ASO problems [3], we performed CFD analysis on wing shape design variations within a certain data space to construct the aerodynamic database and achieved L2 mesh level wing shape optimization by using 140,000 low-fidelity data and 2,000 high-fidelity data. Our previous investigations on multi-fidelity ASO led to the conclusion that an effective approach to reduce reliance on training data is to enhance data quality through efficient geometry sampling.

1.2 Diffusion-based geometry sampling method

During the pursuit of data-driven ASO, a common dilemma arises when striving for high-fidelity, high-dimensional optimization. On the one hand, a large number of aerodynamic data are required for training to ensure accurate surrogate model constructions. However, generating such a large dataset incurs substantial computational costs, which contradicts the objective of using surrogate models in optimization—i.e., to minimize the computational expenses. With this consideration, the quality of the data used to train the surrogate model becomes critical. An efficient geometry sampling method can help the design by enabling the exploration of diverse design alternatives. The general methods are Design of Experiments (DoE) methods that utilize statistical principles to sample within a search space in order to understand the relationship between input variables (factors) and the output response. The representative DoE methods include Latin Hypercube Sampling (LHS) [2], Uniform Design (UD) [24, 25] and Monte Carlo Design (MCD) [26, 27]. With specific constraint settings and non-linear programming, some methods also achieve sampling within constraint handling including probability methods (e.g., constrained Expected Improvement [EI] and Probability of Improvement [PI]) and traditional methods (e.g., penalty function and sequential quadratic programming [SQP]) [5].

Diffusion models offer several advantages over traditional sampling approaches. They excel in searching design spaces that may contain multimodal or high-dimensional landscapes with numerous local minima or maxima. Moreover, diffusion models can effectively handle non-differentiable or discontinuous objective functions. Unlike traditional methods relying on derivatives, diffusion models—which are based on stochastic dynamics—do not require differentiability, making them highly versa-

tile in handling a wider range of objective functions. Additionally, these models can incorporate prior knowledge, enabling customized and guided exploration. The diffusion model has found application as a sampler in various scenarios for addressing inverse-design problems. These include fields such as bio-mechanics [28], aerodynamic design [29], and chemical design [30]. These diverse applications demonstrate the versatility and effectiveness of the diffusion model in generating solutions for complex design challenges across different domains. The ability of capturing complex dependencies and generating realistic responses further expands the applicability of diffusion models to generate more natural conversations and physical visualizations, extending even to fields such as Chat Bot (e.g., ChatGPT ¹) and video generations (e.g., Sora ²).

In this study, we apply the conditional diffusion model as a geometry sampling method to achieve high-fidelity design optimization problem. To achieve a sampling method that can address the limitations of conventional methods, we aim to meet two conditions. Firstly, the desired sampling method should be able to effectively leverage the physical priors of high-performance design as the guidance for sampling. Secondly, this method should effectively define the design space for high-performance designs and generate a minimal amount of sampled data, if subsequent simulations or experimental validations of this data are still required. Details on how the derived algorithms can meet these conditions will be presented in Section 3, and the experiment results will be presented in Section 4.

This paper is organized as follows. In Section 2, we provide a comprehensive overview of the problem, which encompasses the data-driven wing shape design framework, the utilization of a CFD solver to generate training data, and the establishment of multiple fidelity levels of data. In Section 3, we present our proposed methodology the conditional diffusion-based geometry sampling method. In Section 4, we present the experiments and discuss the results from three aspects: the generative aerodynamic data using the diffusion model 4.1, the multi-fidelity design explorations 4.2 and the generative data support data-driven ASO and further operation-aware wing design 4.3. In Section 5, we provide concluding remarks on our method and experiments, while discussing the limitations and potential areas for future research.

2 Problem Description

The main objective of this study is to explore an effective geometry sampling method to reduce the demand for high-fidelity CFD-based aerodynamic data quantity, and further decrease the computational cost involved in establishing data-driven and multi-fidelity ASO. We establish the framework and validate the approach on a single-point ASO problem, using the National Aeronautics and Space Administration (NASA) Common Research Model (CRM) wing configuration as the baseline configuration³. In this section, we provide an overview of the framework software platform, geometry, and data features essential to our study. In Section 2.1, we present the gradient-based platform that forms the foundation for our method. Our approach is built upon this platform, which allows us to leverage gradient information to optimize the design process effectively. We also introduce the shape parameterization method that defines the data features of our aerodynamic data. In Section 2.2, we introduce the CFD solver utilized for analyzing the wing meshes. In Section 2.3, we delve into the training aerodynamic data used for our data-driven design approach, along with the associated data features. The flow of the wing shape design optimization with conditional diffusion-based geometry sampling method is illustrated in Figure 1.

¹<https://chat.openai.com/> (last accessed on 14 June 2024).

²<https://openai.com/index/sora/> (last accessed on 14 June 2024).

³<https://commonresearchmodel.larc.nasa.gov/home-2/high-speed-crm/experimental-approach/model-description/> (last accessed on 14 June 2024).




Figure 1 – Work flow of data-driven wing shape design optimization with conditional diffusion-based geometry sampling.

2.1 Data-driven wing shape design with diffusion-based sampling

As previously investigated in Section 1.1, using gradient-based optimizers—combined with deep NN as surrogate model—offers an appropriate approach to achieve an efficient data-driven ASO. This approach has two main advantages. Firstly, using DNNs as surrogate models allows for efficient computation of gradients through backpropagation, which can then be used as the gradient information and further support the sensitivity analysis of surrogate-based ASO. Secondly, DNNs can effectively provide surrogates for high-dimensional data with high-precision regression. By leveraging gradient information and DNNs, the gradient-based framework can support ASO efficiently.

This surrogate-based (or data-driven) ASO gradient-based framework with DNN as surrogate model has been demonstrated successfully in a previous study by one of the co-authors [2]. In that study, five layers of fully connected artificial neural networks (ANNs) were used as the surrogate model to handle a 60-dimensional ASO problem, with 140,000 aerodynamic data that contain the information $\{M, \alpha, h, \alpha_{\text{twist}}, z\}$ as inputs. The surrogate models produce aerodynamic coefficients $\{C_L, C_D, C_M\}$ as outputs. The diagram shown in Figure 2 illustrates the data-driven ASO framework with single-fidelity aerodynamic data.

We develop our data-driven wing shape design framework based on a gradient-based design optimization framework `MACH-Aero`⁴, which was developed by the MDO Lab at the University of Michigan⁵. In the ASO problem presented in this paper, we use the compact modal parameterization approach developed by Li and Zhang [32] for the wing shape parameterization. The efficiency of ASO with modal parameterization approach has been validated in previous works [33, 34, 35]. As a further extension of this approach, Li and Zhang [32] proposed a deep-learning-based method to generate realistic wing samples in the desired design space, addressing the absence of global wing mode shapes. Compared to conventional modal parameterization approaches, compact modal parameterization utilizing machine learning techniques generates reliable wing samples and mitigates

⁴<https://mdolab-mach-aero.readthedocs-hosted.com/> (last accessed on 14 June 2024).

⁵<https://mdolab.engin.umich.edu/> (last accessed on 14 June 2024).




Figure 2 – The eXtended Design Structure Matrix (XDSM) diagram [31] of the data-based wing design optimization.

the issue of the absence of public geometric information (lack of openly shared details about the shape and structure of wings, hindering analysis). This method has been successfully demonstrated and proven effective in some previous research [33, 36, 3]. All optimizations presented in this paper are performed by using the sequential least squares programming (SLSQP) optimizer, which is assembled in one of the MACH-Aero modules pyOptSparse⁶.




Figure 3 – CRM wing geometry (with eight sectional airfoils) and the first eight wing shape modes extracted using compact modal parameterization method, we use 50 mode coefficients to manipulate the linear combination of these modes.

With conditional diffusion model (to be further elaborated in Section 3), we have the generative wing samples with desired value function setting (also desired aerodynamic performance requests) trained on L3 mesh aerodynamic data. We then evaluate these shape variations with L2 mesh. The volumes of sampled wings generated by diffusion model are manually controllable. With the consideration of CFD computational costs, we only sample 100 wings with the given value function. We further build a surrogate model on these 100 data points with geometry and aerodynamic coefficients combined. Since these wing samples are generated with the priori information from L3 mesh aerodynamic data,

⁶<https://github.com/mdolab/pyoptsparse.git>.(last accessed on 14 June 2024).

we do not need a composite neural network strategy or fine-tuning strategy to leverage both L3 and L2 aerodynamic data. This approach, which is based on existing data and applies certain physical priors while imposing conditions to generate new data, can also be considered as a form of data augmentation.

2.2 CFD solver

We use `ADflow`⁷ as the CFD solver to generate the aerodynamic data given inputs $x := \{x_{\text{con}}, x_{\text{geo}}\}$ and evaluate the corresponding aerodynamic coefficients $\{C_L, C_D, C_M\}$. ADflow computes and outputs the state variables in physical fields, objective function, and constraint function, which are subsequently used as inputs for the adjoint computation [37]. It solves the Reynolds-averaged Navier-Stokes (RANS) equations using multi-block meshes, in combination with alternative numerical schemes such as the Runge-Kutta, diagonally dominant alternating direction implicit (DDADI), and Newton-Krylov methods, as described in the work by Kenway *et al.* [38]. In particular, we use the approximate Newton-Krylov (ANK) solver [39] in ADflow for fast computation, while Newton-Krylov (NK) is employed to generate the final solutions.

2.3 Multiple fidelity aerodynamic data and data features

There are two types of aerodynamic data involved in this paper. The first type of data is the CFD-based aerodynamic data, which are generated by employing a CFD solver described in Section 2.2 to analyze the wing meshes. We define different fidelity levels of data by conducting CFD simulations using wing meshes of varying fidelity levels. We designate the analysis results of ADflow on the L3 CRM wing mesh (with a mesh size of 56,320) as low-fidelity data, and the analysis results of ADflow on the L2 CRM wing mesh (with a mesh size of 450,560) as high-fidelity data. Figure 4 shows two meshes used in this study. The second set of data contains generative aerodynamic data, which is generated using the diffusion model described in Section 3. This model captures the CFD-based data distribution by simulating a diffusion process in a latent space, allowing for the generation of new samples by reversibly iterating through the diffusion process.




Figure 4 – The low and high-fidelity CRM wing meshes used to generate CFD-based aerodynamic data in this work.

All data in this paper follow the same data feature format. The quantity of interests (QoIs) is $y := \{C_L, C_D, C_M\}$, which are calculated using the CFD solver introduced in Section 2.2. These QoIs are

⁷<https://github.com/mdolab/adflow/> (last accessed on 14 June 2024).

evaluated based on a given set of inputs $x := \{x_{\text{con}}, x_{\text{geo}}\}$, which comprise operating condition information $x_{\text{con}} := \{M, \alpha, h\}$ and geometry information $x_{\text{geo}} := \{\alpha_{\text{twist}}, z\}$, where M is Mach number, α is angle of attack, h is flight altitude, α_{twist} is wing twist angle, and z is wing shape perturbations. As shown in Figure 3, we use compact modal parameterization to extract 50 modes of wing shapes. z is a vector that manipulates the linear combinations of these wing shape modes. By manipulating the various parameters z , we are able to achieve different variations in wing shape.

3 Conditional diffusion-based geometry sampling

In recent years, the diffusion models have demonstrated impressive performance in generation capability [40, 41, 42]. These models not only leverage their advantages in terms of model capacity but also possess the ability to generate new samples in high-dimensional space that adhere to physical laws, exhibit realistic effects, and diverge from the original dataset. Diffusion model-based techniques have achieved good generalization in data sampling. This section is organized as follows. In Section 3.1, we introduce the background of diffusion model. In Section 3.2, we introduce our value-function-guided conditional diffusion model developed for generative design.

3.1 Diffusion model background

The diffusion model is a type of generative model that is used to approximate a target probability distribution by simulating a stochastic process. It is particularly effective in modeling complex high-dimensional data distributions and has shown effectiveness in series of studies [43, 44, 45]. Given a dataset $\mathcal{D} = \{x_i\}_{i=1}^N \sim p_{\text{data}}(x)^{\otimes N}$, where p_{data} concentrates on a feasible set C , we will learn a generative model p such that $p \approx p_{\text{data}}$ using a diffusion model [43, 44, 45]. In our study, x_i is a vector that concatenates flight condition $x_{\text{condition}}$ and geometry design variables x_{geometry} , C represents the design domain space. In diffusion models, we first simulate a forward noising process starting from data distribution $p_0(x) = p_{\text{data}}(x)$ which converges to the standard Gaussian distribution $p_T(x) \approx \mathcal{N}(0, \text{Id})$ as $T \rightarrow \infty$. The forward process is defined by the following stochastic differential equation (SDE) [46],

$$d\mathbf{x}_t^\rightarrow = f(\mathbf{x}_t^\rightarrow, t) dt + g(t) d\mathbf{w}_t, \mathbf{x}_0^\rightarrow \sim p_{\text{data}}(x), 0 \leq t \leq T, \quad (1)$$

where $f: \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a vector-valued drift function, $g(t): \mathbb{R} \rightarrow \mathbb{R}$ is a scalar-valued diffusion coefficient, and $(\mathbf{w}_t)_{t \geq 0}$ is a d -dimensional Brownian motion [47]. Next, with mild assumptions [48, 49] that are satisfied for the processes in this study, the reverse process that generates the data from normal noise admits a SDE description,

$$d\mathbf{x}_t^\leftarrow = [-f(\mathbf{x}_t^\leftarrow, \tau) + g^2(\tau) \nabla_x \log p_\tau(\mathbf{x}_t^\leftarrow)] dt + g(\tau) d\mathbf{w}_t, \quad (2)$$

where $\tau = T - t$ and $\nabla_x \log p_\tau(x)$ is the score function which is modeled by a time-dependent NN with the score matching objective

$$\mathbb{E}_t \left[\lambda(t) \mathbb{E}_{\mathbf{x}_0^\rightarrow} \mathbb{E}_{\mathbf{x}_t^\rightarrow | \mathbf{x}_0^\rightarrow} \left[\left\| s_\theta(\mathbf{x}_t^\rightarrow, t) - \nabla_x \log p_{t|0}(\mathbf{x}_t^\rightarrow | \mathbf{x}_0^\rightarrow) \right\|_2^2 \right] \right]. \quad (3)$$

In Equation 3, $p_{t|0}(\mathbf{x}_t^\rightarrow | \mathbf{x}_0^\rightarrow)$ is the conditional density of the forward SDE starting at $\mathbf{x}_t^\rightarrow$, and $\lambda(t) > 0$ is a weighting function. The drift $f(t)$ and diffusion $g(\cdot)$ function can express the state of the forward SDE as a linear combination of the initial state and a scaled normal random variable under appropriate assumptions on drift $f(t)$ and diffusion $g(\cdot)$ function [50]. The above process forms a forward and reverse process of a diffusion model.

3.2 Value function-guided conditional diffusion model

As we introduced in Section 3.1, a diffusion model can be used to model the probability distribution of the dataset $\mathcal{D} = \{x_i\}_{i=1}^N \sim p_{\text{data}}(x)^{\otimes N}$. However, in most cases, we require diffusion model to approximate the probability distribution of data samples under a specific condition c , where $\mathcal{D} = \{x_i\}_{i=1}^M \sim p_{\text{data}}(x|c)^{\otimes M}$ —we call this the conditional diffusion model. The type of this condition can be encoded as classification tags [51]. This type of condition could also be an implicit function, which we aim to sample from the following distribution:

$$p(x|c) \propto q(x)e^{-\beta \mathcal{E}(x)}. \quad (4)$$

where $q(x)$ is an unknown data distribution $\mathcal{X} \subseteq \mathbb{R}^d$, $\mathcal{E}(\cdot)$ is the energy function from $\mathcal{X} \subseteq \mathbb{R}^d$ to $\mathbb{R}$. The energy function $\mathcal{E}(\cdot)$ can be chosen flexibly, as long as the integral of $q(x)e^{-\beta \mathcal{E}(x)}$ over $x \in \mathcal{X}$ is preserved. Readers are referred to the work by Lu *et al.* [52] for further details of the definition and theoretical background of energy function.

In this study, we propose a value function as the guidance energy function for the diffusion model. There are three consecutive processes to achieve this maximum value function-guided diffusion model. Each process contains a loss function as the training objective of the model. The three processes have different responsibilities for the overall model functionality with the associated loss function also serves different training objectives. We take these processes as three independent models. The first diffusion model, denoted as $\mu(a|s)$, learns the distribution of the given data. Here, a represents the sampling design variables and s represents the conditions. Additionally, we utilize this diffusion model to model the desired performance policy (specifically aerodynamic performance in this study), denoted as $\varepsilon_\theta(a_t|s, t)$, where t represents the time state of the diffusion model and θ represents the model parameters. The second model is a maximum value function that depicts the condition we impose on the diffusion model. In this study, the value function contains the desired aerodynamic performance information we expect from generative model and the information is quantified with aerodynamic coefficients C_L , C_D and C_M . The third model is an energy model $f_\phi(a_t, s, t)$, which is used to evaluate the intermediate state of diffusion process and also guide the maximum value function conditional diffusion model introduced from previous two processes. We will introduce these three models in details in this section and the notations will be illustrated in the subsequent paragraphs.

A diffusion model to learn the given data distributions

Firstly we build a diffusion model training on a given aerodynamic dataset $\mathcal{D}_\mu$. Referring to the theory from score-based generative modeling through stochastic differential equations [46], we introduce Equation 5 as the training loss function,

$$\begin{aligned} \theta^* &= \arg \min_{\theta} \mathbb{E}_{(c,a) \sim \mathcal{D}_\mu, \varepsilon, t} \left[\|\sigma_t \nabla_{a_t} \log \mu_t(a_t|c) + \varepsilon\|_2^2 \right], \\ &= \arg \min_{\theta} \mathbb{E}_{(c,a) \sim \mathcal{D}_\mu, \varepsilon, t} \left[\|\varepsilon_\theta(a_t|c, t) - \varepsilon\|_2^2 \right], \end{aligned} \quad (5)$$

where θ^* represents the first process diffusion model parameters, $\varepsilon_\theta(\cdot)$ represents the noise predicted by the model with parameters θ . Generally, diffusion model introduces the noise-weighted score function to compute the loss function in order to stabilize the training process. In Equation 5, the objective is to minimize the error term, which is the squared Euclidean distance between the predicted noise $\varepsilon_\theta(a_t|c, t)$ and the actual noise ε . Here, a_t represents the sampled design variables at a specific time step t , and c represents the condition. In our model, the diffusion model with parameter θ predicts the noise $\varepsilon_\theta(a_t|c, t)$ for the current design, which is used to scale the design perturbed by diffusion process. Given the condition c and time step t , we can express the design perturbed by diffusion process following a_t as $a_t = \alpha_t a + \sigma_t \varepsilon$, where α_t is a scaling factor, a represents the original design variables, σ_t represents the diffusion noise scale, and ε is a sample from the normal

distribution $\varepsilon \sim \mathcal{N}(0, I)$. To scale the noise predicted by the diffusion model to the desired normal distribution, we have the relationship:

$$\varepsilon_\theta(a_t | c, t) = -\sigma_t \nabla_{a_t} \log \mu_t(a_t | c). \quad (6)$$

Here, $\mu_t(a_t | c)$ represents the diffusion model's prediction of the probability distribution of the diffusion-augmented action a_t given the design variables c . The negative gradient of the logarithm of this probability distribution is multiplied by the diffusion noise scale σ_t to obtain the noise term $\varepsilon_\theta(a_t | c, t)$. In this procedure, the diffusion model predicts the noise for the current design, and to scale this noise to a normal distribution, we use the relationship 6, where the design perturbed by diffusion process a_t is obtained by adding the scaled diffusion noise to the original design variables a .

Maximum value function

The maximum value function is a condition imposed on the diffusion model. This function should be designed as a desired performance that we expect from the generative data from the diffusion model. Specifically in this study, we use diffusion model as a sampler to generate the wing geometry data with high aerodynamic performance. Thus, the maximum value function is designed to contain the desired aerodynamic evaluations information. For instance, we would like to generate the wing geometry that has sufficient lift and stability under the aerodynamic conditions of Mach = 0.85 and altitude = 11,740 m, referring to the setting from the American Institute of Aeronautics and Astronautics (AIAA) Aerodynamic Design Optimization Discussion Group (ADODG)⁸ case 4.5. The desired wing should maintain $C_L = 0.5$, with $C_M > -0.17$, and with C_D as small as possible, which means the aircraft drag reduction characteristics is better with evaluated C_D is smaller. Based on these requirements, we design our value function V as shown in Equation 7:

$$V = \beta \left[-(\lambda_1 C_D - 0.018)^2 - \lambda_2 (C_L - 0.5)^2 - \lambda_3 \max(C_M, -0.17)^2 \right], \quad (7)$$

where β represents the sampling intensity with respect to the overall data distribution. A higher β value indicates that we are sampling from a narrower space, leading to generated data that is closer to our desired conditions. Following to the definition of temperature coefficient in the energy equation [53], we also call β the temperature coefficient. $\lambda_{1,2,3}$ are the hyperparameters, we tune these parameters to control the relative importance of each sub-condition depicting C_D , C_L and C_M . For future research or to serve other desired aerodynamic performance wing designs, this value function needs to be specifically defined based on the particular problem.

The intermediate state energy function

Suppose we have a well-defined value function that is appropriate for investigating the problem at hand. In this case, we can utilize the Contrastive Energy Prediction (CEP) method [54] to train an intermediate state energy equation that is suitable for the current guidance policy. For a K step training process, the training objective is

$$\arg \min_{\phi} \mathbb{E}_{t,c,\varepsilon^{1:K}, a^{1:K} \sim \mu(a|c)} \left[-\sum_{i=1}^K \frac{e^{V(a^i, c)}}{\sum_{j=1}^K e^{V(a^j, c)}} \log \frac{e^{-f_\phi(a^i, c, t)}}{\sum_{j=1}^K e^{-f_\phi(a^j, c, t)}} \right], \quad (8)$$

where $\varepsilon^{1:K}$ is a sequence of K noise samples along with K training steps, $a^{1:K} \sim \mu(a | c)$ is a sequence of K actions, where each perturbed design variables a are sampled from the policy $\mu(a | c)$.

After completing the training of the intermediate state energy function, we combine it with the value function as an optimal policy to guide the generation of diffusion model. The Algorithm 1 depicts

⁸<https://sites.google.com/view/mcgill-computational-aerogroup/adodg> (last accessed on 14 June 2024).

the above whole process of achieving value function guided conditional diffusion model designed for generating high aerodynamic performance wings. The value function $V(\cdot)$ here is the energy function $\mathcal{E}$.

Algorithm 1 Diffusion Model as a Sampler for Generative Design

```

1: Input: Given CFD analyzed training wing dataset  $\mathcal{D}_\mu$ , which contains the design variable  $a$  and
   condition  $c$  information.
2: //Training the diffusion model
3: for each gradient step do:
4:   Sample  $B$  data points  $(a, c)$  from  $\mathcal{D}^\mu$ ,  $B$  Gaussian noises  $\varepsilon$  from  $\mathcal{N}(0, I)$ , and  $B$  time steps  $t$ 
   from  $\mathcal{U}(0, T)$ 
5:   Perturb  $a$  according to  $a_t := \alpha_t a + \sigma_t \varepsilon$ 
6:   Update  $\theta \leftarrow \theta - \lambda_\theta \nabla_\theta \sum [\|\varepsilon_\theta(a_t | s, t) - \varepsilon\|_2^2]$ 
7: end for
8: //Generating perturbed designs using the diffusion model
9: for each state  $c$  in  $\mathcal{D}^\mu$  do:
10:  Sample  $K$  perturbed designs  $\hat{a}^{(1:K)}$  from the diffusion model  $\mu_\theta(\cdot | c)$  and store them as  $\mathcal{D}^{\mu_\theta}(c)$ 
11: end for
12: //Training the evaluation model and the energy guidance model
13: for each gradient step do:
14:  Sample  $B$  data points  $(c, a, r, c')$  from  $\mathcal{D}^\mu$ ,  $B$  Gaussian noises  $\varepsilon$  from  $\mathcal{N}(0, I)$ , and  $B$  time steps
    $t$  from  $\mathcal{U}(0, T)$ 
15:  Calculate the target Value function  $\mathcal{T}^\pi V_\psi(c, a)$  and detach gradient
16:  Update  $\psi \leftarrow \psi - \lambda_\psi \nabla_\psi \sum [\|V_\psi(c, a) - \mathcal{T}^\pi V_\psi(c, a)\|_2^2]$ 
17:  Perturb  $\hat{a}$  according to  $\hat{a}_t := \alpha_t \hat{a} + \sigma_t \varepsilon$ 
18:  Update  $\phi \leftarrow \phi + \lambda_\phi \nabla_\phi \sum_i \left[ \frac{e^{\beta V_\psi(c, \hat{a}_i)}}{\sum_j e^{\beta V_\psi(c, \hat{a}_j)}} \log \frac{e^{f_\phi(\hat{a}_{i,t}|c,t)}}{\sum_j e^{f_\phi(\hat{a}_{j,t}|c,t)}} \right]$ 
19: end for
20: Output: Generated wing samples  $\mathcal{D}_\mu$ 

```

4 Experiments and Discussions

In this section, we present the results and discussions of our study. The purpose of the experiment is to validate whether the proposed diffusion-based geometry sampling method can support the data-driven ASO, multi-fidelity design exploration and further operation-aware design facing rapid industry design demand. To support this argument, we will discuss in following three sections. Firstly in Section 4.1, we will compare the aerodynamic performance of the wings sampled using our diffusion-based method with those sampled using the LHS method. Secondly in Section 4.2, we will make the high-fidelity design explorations with low-fidelity data. Thirdly in Section 4.3, we will perform some data-driven design optimization cases using the diffusion-based generative data.

4.1 Generative wing samples with conditional diffusion model

In this section, we present wing data generation results using conditional diffusion-based sampling method. We aim to experimentally verify and address two questions. First, does the conditional diffusion-based sampling method generative geometry data satisfy the given conditions? Second, to what extent does the performance of the generated geometry from the conditional diffusion-based sampling method decrease as the mesh fidelity level increases? In other words, if we generate high L/D geometry using conditional diffusion-based sampling method, will it still maintain a high L/D ratio

as the fidelity level increases? If there is a decrease, what is the quantification of this reduction? To answer this questions, we do experiments as follows.

We follow the conditional diffusion-based sampling procedure described in Section 3 and use the value function as conditions described in Equation 7. We use 135,108 L3 CRM training data with the same data features as depicted in Section 2.3. These datasets are generated by applying LHS perturbations to the CRM baseline within a specified design space. We utilize `IDWarp`⁹ for mesh warping and `ADFlow`¹⁰ as the CFD solver to analyze the perturbed wing meshes and obtain C_L , C_D , C_M evaluations. These datasets have also shown effectiveness to support data-driven ASO of CRM wing on solving multiple design optimization problems in previous investigation [2]. We will employ this dataset as the training set to train the diffusion model. Subsequently, we use the CFD solver to evaluate and calculate the L/D ratios of the generated wings. Additionally, we adjust the alignment of the generated outputs with the target conditions by configuring different temperature coefficients β in the value function 7.

The specific definition and mathematical expression for the value function (Equation 7 in Section 3)—which applies constraints to the wing sample and provides conditional guidance for the diffusion model—depend on the specific optimization problem at hand. For the optimization problem presented in this paper, the sampled wings need to satisfy that $C_L = 0.5$, $C_M > -0.17$, and $C_D \approx 0.0229$ (evaluated at Mach number 0.85 and flight altitude 11,740 m). The value function is then formulated as

$$\begin{aligned} V(x) = & \beta \times [-(C_{Dest} * 10 - 0.023)^2 \\ & - \lambda_1 (\max(\max(C_{Lest} - 0.505, 0), \max(0, 0.495 - C_{Lest})))^2 \\ & - \lambda_2 (\max(\max(C_{Mest} + 0.168, 0), \max(0, C_{Mest} + 0.170)))^2. \end{aligned} \quad (9)$$

It is worth noting that the condition $C_D \approx 0.0229$ is derived *a priori* through CFD-based optimization. From the consideration of experimental efficiency, we utilize this prior information as part of the value function information. To increase the variations in the sample data and prepare for higher-fidelity data in the sequential process, we allocate 1% margin for each coefficient in the value function. This is because high-fidelity aerodynamic coefficient evaluations differ significantly from low-fidelity evaluations, and we do not want to restrict the design space strictly during the initial sampling process. To ensure consistent thresholds for the variations of C_L , C_D , and C_M , we magnify the variations of the C_D value by a factor of 10. This adjustment allows us to align balanced thresholds of these parameters. When applying this approach to future optimization problems, the value function should be designed specifically based on their specific requirements.

We use the value function shown in Equation 9 and follow the conditional diffusion-based geometry sampling process introduced in Section 3 to perform the experiments. The temperature coefficient β affects the *sharpness* of the generated samples. A higher temperature value produces sharper, more confident samples, while a lower temperature value results in more diffuse, less confident samples. By varying the temperature coefficients β , we select different sharpness of wing samples. The sampling results are summarized in Table 1. We make a statistical analysis about different distribution of wing samples across different L/D ranges for each beta setting. Furthermore, we compare and visualize the diffusion-based sampling data distributions to LHS data distributions (specifically contains 140,000 aerodynamic data) shown in Figure 6.

From Table 1 and Figure 6, we can see that that compared to the LHS-based method, the diffusion-based sampling generative data are more concentrated on higher L/D ratio area. Following the theory of diffusion model, the data distributions of diffusion-based sampling follows the Gaussian distribution. As β increases, the diffusion-based wing samples data distributions become sharper. An interesting observation is that we can hypothesize that as β approaches infinity, the data distribution

⁹<https://github.com/mdolab/idwarp.git> (last accessed on 14 June 2024).

¹⁰<https://github.com/mdolab/adflow/> (last accessed on 14 June 2024).

L/D threshold	Proportion				
	$\beta = 8$	$\beta = 16$	$\beta = 32$	$\beta = 64$	$\beta = 128$
≤ 18	0.08	0.12	0.03	0	0.06
18 – 19	0.27	0.23	0.16	0.07	0.16
19 – 20	0.41	0.34	0.45	0.48	0.51
≥ 20	0.24	0.31	0.36	0.45	0.27

Table 1 – Generative wing samples L/D threshold with varying temperature coefficient β .




Figure 5 – The box plot and probability density plot of conditional diffusion-based sampling method generates better aerodynamic performance (higher L/D) data than LHS method.

will converge towards the mean of Gaussian distribution, resembling a Dirac distribution [55]. In such a scenario, if we set the sample quantity to 1, it raises the question of whether we can obtain an optimal solution. This is a topic worthy of discussion in future work.

We also visualize part of the generative wing samples from diffusion-based sampling method shown in Figure 9. The visualized C_p and shock analysis of wing samples demonstrate the generative wing samples covers enough design space and complex aerodynamic phenomena that support data-driven design optimization.

4.2 Generative wing samples for multi-fidelity design explorations

In this section, we would like to make the multi-fidelity design explorations using conditional diffusion-based geometry sampling method. Using low-fidelity data (specifically L3 mesh level aerodynamic data), we want to sample the wings with higher L/D ratio by adjusting the associated value functions of conditional diffusion model. Specifically, we maintain $C_L = 0.5$ while reducing C_D , resulting in an increased L/D ratio. In fact, achieving the desired C_D values proves challenging on low-fidelity training data using traditional LHS method. Therefore, we seek to explore whether the diffusion model can achieve these high aerodynamic performance (high L/D ratio) wing samples by learning




Figure 6 – The conditional diffusion-based sampling generated data distributions varying by different temperature coefficient β .

from low-fidelity aerodynamic data that traditional sampling method cannot achieve. Specifically, we aim to determine if the model can explore design variations that maintain high performance even when subjected to high-fidelity CFD validation on a finer mesh level. Referring to the Equation 9, we adjust the value function by reducing the C_D from 0.023 to 0.019. Based on the experiments shown in Section 4.1, we use temperature coefficient $\beta = 64$ as the diffusion model guidance scale value. Upon obtaining the generated wing geometry, we utilize the mesh wrapper `IDWarp`, and the CFD solver `ADflow`, to estimate the aerodynamic coefficients at the L2 mesh level for these wing samples.

Specifically, the generative geometry data still follow the same features as depicted in Section 2.3, which is $\{M, \alpha, h, \alpha_{\text{twist}}, z\}$. We then analyze these sampled wings using `ADflow` at L2 and L3 meshes. We analyze the C_L, C_D, C_M for these wings and calculate the corresponding L/D ratios. We generate probability distribution plots for C_L, C_D, C_M , and L/D , under CFD evaluation for both L2 and L3 meshes. The data distributions for L2 and L3 mesh level diffusion-based generative aerodynamic data are shown in Figure 8. From the Figure 8, we compare three different cases of sampling. The notations for each sampling case are as follows. "L3 Data - LHS": Geometry sampled using the LHS method and solved using CFD on the L3 mesh; "L2 Data - LHS": Geometry sampled using the LHS method and solved using CFD on the L2 mesh. "Diffusion": Geometry sampled using the conditional diffusion-based sampling method and solved using CFD on the L2 mesh. The data volume of each case is: "L3 Data - LHS": 140,000, "L2 Data - LHS": 2,000, "Diffusion": 500. From this figure, it is evident that the diffusion model is capable of generating wing samples with high L/D ratios that could not have been achieved by using the LHS method. This provides us with a greater number of high quality design samples and design explorations for future multi-fidelity design optimization.

4.3 Generative data support data-driven wing shape design optimization

In this section, we present how conditional diffusion-based generative data support data-driven ASO. The procedures of data-driven ASO experiments are as follows: we build the surrogate model of our data-driven optimization framework based on the given training aerodynamic data. The trained




Figure 7 – CFD evaluations on generative wing samples, C_p distributions and shock analysis indicate that the design space contains complex aerodynamic phenomena.

surrogate model behaves as the aerodynamic coefficient estimator in the data-driven optimization framework. We use multilayer perceptron (MLP) as the basic structure of surrogate model to perform data-driven ASO using only one fidelity level (specifically L3 level of CRM mesh in this experiment) of aerodynamic data. We train the surrogate models using data generated from different sampling methods, LHS and our conditional diffusion-based sampling method described in Section 3. These methods result in varying quantities of training data for training the surrogate models. The main objective of this experiment is to validate the impact of two different sampling methods on data-driven ASO. In order to ensure a fair comparison, we use the MLP as the basic framework for training, rather than other forms of neural networks. The specific configurations, such as the depth and number of nodes in the network, varies depending on the size of the training dataset. We use forward differentiation to estimate the gradient of the surrogate model when utilizing it for gradient-based optimization. This method serves as an approximation technique to compute the gradients of the objective function or constraints with respect to the input variables. Afterwards, we perform the data-driven ASO under a gradient-based optimization framework using a nonlinear programming optimizer *SLSQP*, which is integrated in *pyOptSparse*¹¹. After completing the data-driven procedure, we proceed to analyze the optimized wing using a CFD solver for validation. We then compare the results of the data-driven optimization with those obtained using a conventional approach, which relies on CFD for estimating aerodynamic coefficients.

We verify this data-driven optimization framework in a drag minimization problem summarized in Table 2, following the guidance of the ADODG case 4.1¹². The operating conditions of wing are set to flight altitude 11,740 m and Mach number 0.85. We employ strict thickness constraints rather than a loose volume constraint to ensure practical design [33]. As stated in Section 2.1, we utilize a compact modal parameterization method for wing geometry parameterization. Consequently, we employ wing shape mode coefficients to govern the perturbations in wing shape, which are also designated as the

¹¹<https://github.com/mdolab/pyoptsparse.git> (last accessed on 14 June 2024).

¹²<https://sites.google.com/view/mcgill-computational-aerogroup/adodg> (last accessed on 14 June 2024).




Figure 8 – The box plot and probability density plot show that diffusion-based sampling method generates better aerodynamic performance (larger L/D) data validated on L2 mesh than directly LHS on L2 mesh data.

design variables. For the generated data with conditional diffusion model, we leverage fine-tuning techniques to modify the output layer of a trained surrogate model using 70,000 LHS dataset. Fine-tuning [56, 57] is a widely utilized technique in meta models, aiming to improve the model’s accuracy on specific data while preserving its generalization capabilities.

The running results are summarized in Table 3. For the LHS, our investigations on training data are mainly on training data volume, data selection percentage; For the diffusion-based sampling method, our investigations on training data are on training data volume, temperature coefficient β and calculated L/D term in value function setting. We define a successful data-driven ASO as the optimization of the wing shape that, when evaluated using CFD solver, satisfies the constraint requirements of the optimization problem 2. We calculate the relative errors ε of aerodynamic coefficients estimations using trained surrogate models,

$$\varepsilon = |(y_{SM} - y_{CFD}) / y_{CFD}| \times 100\%, \quad (10)$$

where y_{SM} and y_{CFD} represent the estimation value of corresponding aerodynamic coefficients obtained using surrogate model and CFD simulations, respectively.

From the experiments summarized in Table 3, we can observe that compared to the LHS method, diffusion-based method generates higher L/D ratio wing samples based on the given value function. The results obtained exclude the parts of wing samples that have low L/D ratio. While LHS offers a comprehensive search space for design, it also introduces a substantial number of low aerodynamic performance wing samples. This increases the computational burden for subsequent CFD calculations on these samples. In contrast, the conditional diffusion model generates wing samples guided by constraints on parameters such as C_L , C_D , and C_M . As a result, the diffusion-based sampling method produces data with higher L/D ratios. Additionally, since the diffusion model is trained using aerodynamic data generated from CFD simulations, the sampling process aligns with this physics-based prior knowledge, ensuring the physical validity of the generated wing samples.

	Function/variable	Description	Quantity
Minimize	C_D	Drag coefficient	1
With respect to	α	Angle of attack	1
	α_{twist}	Angles of twists	7
	z	Wing shape modes perturbation	50
		Total design variables	58
Subject to	$C_L = 0.5$	Lift constraint	1
	$C_{M_y} \geq -0.17$	Pitching moment constraint	1
	$t \geq 0.98 \times t_{\text{initial}}$	Thickness constraints	750
	$V \leq 1.02 \times V_{\text{initial}}$	Maximum volume constraint	1
	$V \geq V_{\text{initial}}$	Minimum volume constraint	1
	$\Delta z_{\text{LE, upper, root}} = -\Delta z_{\text{LE, lower, root}}$	Fixed leading-edge constraints	8
	$\Delta z_{\text{TE, upper}} = -\Delta z_{\text{TE, lower}}$	Fixed trailing-edge constraints	8
		Total constraints	770

Table 2 – ASO problem formulation as a drag minimization problem.

		LHS				Diffusion-based sampling	
Training data	Training data volume	130,000	70,000	10,000	500	500	500
	Percentage (%) / β	100%	50%	7.6%	0.38%	32	64
	Value function (L/D)	-	-	-	-	21	20
	L/D bounds	9.45 – 22.66				15.56 – 22.6	18.17 – 22.06
Surrogate model	$\varepsilon(C_L)$	0.20	0.39	1.10	3.20	0.15	0.14
	$\varepsilon(C_D)$	0.35	0.62	1.94	5.44	0.24	0.22
	$\varepsilon(C_M)$	0.36	0.70	1.77	4.93	0.29	0.27
	Succeed	✓	-	-	-	✓	✓
Optimization	C_L	0.497	0.5	0.501	0.501	0.499	0.499
	C_D (counts)	229	267	261	279	229	229
	C_M	-0.17	-0.205	-0.2	-0.22	-0.17	-0.17
Computational cost	Core hours	56,000	28,000	4,000	200	200	200

Table 3 – Data-driven ASO using LHS and diffusion-based sampling method with different size of training data.

5 Conclusion

In this paper, we presented an effective conditional diffusion-based geometry sampling method that could generate high aerodynamic performance data with associated value function designed. Inspired by the energy conditional diffusion model theory, we proposed a value function-guided diffusion model to generate high performance wing samples to offer an efficient mean to yield high-fidelity wing design. We validated the effectiveness of this conditional diffusion-based sampling method in three key aspects, namely in generating desired high L/D ratio wing samples, supporting multi-fidelity design explorations, and enabling data-driven wing shape design optimization. The associated experiments have shown that with appropriate value function designed, the conditional diffusion model can generate desired L/D ratio wing samples with a normal distribution. Furthermore, the wing samples trained using low-fidelity data continue to adhere to this distribution when evaluated on a high-fidelity mesh, satisfying our requirements for the aerodynamic performance of wing samples under high-fidelity design. We fine-tuned our model to enable this generative data to support our data-driven aerodynamic shape optimization and get the results obtained are competitive with CFD-based design optimization.

Specifically, in each experiment, the experiment in Section 4.1 shows that with desired value function designed, we can generate the desired $L/D = 20$ wing samples using conditional diffusion model. In Section 4.2, we validate that we can generate higher L/D ratio wing samples from low-fidelity data that fits for a higher fidelity design. And with value function adjustment, we successfully generate the wing samples that follow a Gaussian distributions with a mean value around $L/D = 23$. In Section 4.3, we successfully use the generated wing samples to fine-tuning a preliminary surrogate model trained with LHS data. By adopting this approach, we can reduce the demand for the training set by 50.5%, while also decreasing the required CFD computations. It is worth noting that the establishment of this training set is aimed at addressing a series of optimization designs, as well as meeting the rapid high-fidelity design requirements for future mission-based and operation-aware aircraft designs.

In future work, we will investigate more geometry constraints imposed on the diffusion model (including wing span area, wing thickness, aspect ratio, etc.) to support more comprehensive design explorations. In the meantime, we aim to extend the exploration of data not only limited to high-fidelity design but also to include a broader range of wing structures for various types of aircraft and flight missions. This expansion will allow us to tackle design tasks that cater to different aircraft and flight mission requirements, thereby supporting an efficient operation-aware aircraft design optimization.

Acknowledgments

This work was supported by the Hong Kong Research Grant Council General Research Fund (Project No. 16207523). The computational fluid dynamics simulations were performed on the Beijing Paratera supercomputing platform.

6 Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.

References

- [1] Steven L Brunton, J Nathan Kutz, Krithika Manohar, Aleksandr Y Aravkin, Kristi Morgansen, Jennifer Klemisch, Nicholas Goebel, James Buttrick, Jeffrey Poskin, Adriana W Blom-Schieber, et al. Data-driven aerospace engineering: reframing the industry with machine learning. *AIAA Journal*, 59(8):2820–2847, 2021.
- [2] Jichao Li and Mengqi Zhang. Data-based approach for wing shape design optimization. *Aerospace Science and Technology*, 112:106639, 2021.
- [3] Aobo Yang, Jichao Li, and Rhea P Liem. Multi-fidelity data-driven aerodynamic shape optimization of wings with composite neural networks. In *AIAA AVIATION 2023 Forum*, page 3470, 2023.
- [4] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 56(4):1–39, 2023.
- [5] Zhong-Hua Han. Surroopt: a generic surrogate-based optimization code for aerodynamic and multidisciplinary design. In *Proceedings of ICAS*, volume 2016, pages 2016–0281, 2016.
- [6] Joaquim RRA Martins and Andrew Ning. *Engineering design optimization*. Cambridge University Press, 2021.
- [7] Luis Miguel Rios and Nikolaos V Sahinidis. Derivative-free optimization: a review of algorithms and comparison of software implementations. *Journal of Global Optimization*, 56(3):1247–1293, 2013.
- [8] Jichao Li and Jinsheng Cai. Massively multipoint aerodynamic shape design via surrogate-assisted gradient-based optimization. *AIAA Journal*, 58(5):1949–1963, 2020.
- [9] Zhonghua Han, Keshi Zhang, Wenping Song, and Jun Liu. Surrogate-based aerodynamic shape optimization with application to wind turbine airfoils. In *51st AIAA aerospace sciences meeting including the new horizons forum and aerospace exposition*, page 1108, 2013.
- [10] Dev Rajnarayan, Alex Haas, and Ilan Kroo. A multifidelity gradient-free optimization method and application to aerodynamic design. In *12th AIAA/ISSMO multidisciplinary analysis and optimization conference*, page 6020, 2008.
- [11] Emiliano Iuliano. Global optimization of benchmark aerodynamic cases using physics-based surrogate models. *Aerospace Science and Technology*, 67:273–286, 2017.
- [12] Akira Oyama. Wing design using evolutionary algorithms. *Department of Aeronautics and Space Engineering of Tohoku University*, 2000.
- [13] Shinichi Takahashi, Shigeru Obayashi, and Kazuhiro Nakahashi. Inverse design optimization of transonic wings based on multi-objective genetic algorithms. *AIAA journal*, 37(12):1656–1662, 1999.
- [14] Imran Ali Chaudhry and Ali Ahmed. Preliminary aircraft design optimization using genetic algorithms. 2014.
- [15] Timothy MS Jim, Ghifari A Faza, Pramudita S Palar, and Koji Shimoyama. Bayesian optimization of a low-boom supersonic wing planform. *AIAA journal*, 59(11):4514–4529, 2021.
- [16] Remy Priem, Hugo Gagnon, Ian Chittick, Stephane Dufresne, Youssef Diouane, and Nathalie Bartoli. An efficient application of bayesian optimization to an industrial mdo framework for aircraft design. In *AIAA Aviation 2020 Forum*, page 3152, 2020.
- [17] Paul Saves, Nathalie Bartoli, Youssef Diouane, Thierry Lefebvre, Joseph Morlier, Christophe David, Eric Nguyen Van, and Sébastien Defoort. Bayesian optimization for mixed variables using an adaptive dimension reduction process: applications to aircraft design. In *AIAA SCITECH 2022 Forum*, page 0082, 2022.
- [18] Loïc Brevault, Mathieu Balesdent, and Ali Hebbal. Overview of gaussian process based multi-fidelity techniques with variable relationship between fidelities, application to aerospace systems. *Aerospace Science and Technology*, 107:106339, 2020.
- [19] Marc C Kennedy and Anthony O'Hagan. Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1):1–13, 2000.
- [20] George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.
- [21] Xuhui Meng and George Em Karniadakis. A composite neural network that learns from multi-fidelity data: Application to function approximation and inverse pde problems. *Journal of Computational Physics*, 401:109020, 2020.
- [22] Mengwu Guo, Andrea Manzoni, Maurice Amendt, Paolo Conti, and Jan S Hesthaven. Multi-fidelity regression using artificial neural networks: efficient approximation of parameter-dependent output quantities. *Computer methods in applied mechanics and engineering*, 389:114378, 2022.

- [23] Xinshuai Zhang, Fangfang Xie, Tingwei Ji, Zaoxu Zhu, and Yao Zheng. Multi-fidelity deep neural network surrogate model for aerodynamic shape optimization. *Computer Methods in Applied Mechanics and Engineering*, 373:113485, 2021.
- [24] Zhiguang Qian, Carolyn Conner Seepersad, V Roshan Joseph, Janet K Allen, and CF Jeff Wu. Building surrogate models based on detailed and approximate simulations. 2006.
- [25] Zhong-Hua Han, Yu Zhang, Chen-Xing Song, and Ke-Shi Zhang. Weighted gradient-enhanced kriging for high-dimensional surrogate modeling and design optimization. *Aiaa Journal*, 55(12):4330–4346, 2017.
- [26] Muhammad Shahbaz, Zhong-Hua Han, WP Song, and M Nadeem Aizud. Surrogate-based robust design optimization of airfoil using inexpensive monte carlo method. In *2016 13th International Bhurban Conference on Applied Sciences and Technology (IBCAST)*, pages 497–504. IEEE, 2016.
- [27] Barron J Bichon, Michael S Eldred, Sankaran Mahadevan, and John M McFarland. Efficient global surrogate modeling for reliability-based design optimization. *Journal of Mechanical Design*, 135(1):011009, 2013.
- [28] Hyungjin Chung, Suhyeon Lee, and Jong Chul Ye. Decomposed diffusion sampler for accelerating large-scale inverse problems.
- [29] Tailin Wu, Takashi Maruyama, Long Wei, Tao Zhang, Yilun Du, Gianluca Iaccarino, and Jure Leskovec. Compositional generative inverse design. *arXiv preprint arXiv:2401.13171*, 2024.
- [30] Rachel K Luu, Marcin Wysokowski, and Markus J Buehler. Generative discovery of de novo chemical designs using diffusion modeling and transformer deep neural networks with application to deep eutectic solvents. *Applied Physics Letters*, 122(23), 2023.
- [31] Andrew B. Lambe and Joaquim R. R. A. Martins. Extensions to the design structure matrix for the description of multidisciplinary design, analysis, and optimization processes. *Structural and Multidisciplinary Optimization*, 46(2):273–284, August 2012.
- [32] Jichao Li and Mengqi Zhang. Adjoint-free aerodynamic shape optimization of the common research model wing. *AIAA Journal*, 59(6):1990–2000, 2021.
- [33] Jichao Li, Mohamed Amine Bouhlel, and Joaquim RRA Martins. Data-based approach for fast airfoil analysis and optimization. *AIAA Journal*, 57(2):581–596, 2019.
- [34] Daniel J Poole, Christian B Allen, and Thomas CS Rendall. Metric-based mathematical derivation of efficient airfoil design variables. *AIAA Journal*, 53(5):1349–1361, 2015.
- [35] Christian B Allen, Daniel J Poole, and Thomas CS Rendall. Wing aerodynamic optimization using efficient mathematically-extracted modal design variables. *Optimization and Engineering*, 19:453–477, 2018.
- [36] Jichao Li, Mengqi Zhang, Joaquim RRA Martins, and Chang Shu. Efficient aerodynamic shape optimization with deep-learning-based geometric filtering. *AIAA journal*, 58(10):4243–4259, 2020.
- [37] Gaetan K. W. Kenway, Charles A. Mader, Ping He, and Joaquim R. R. A. Martins. Effective adjoint approaches for computational fluid dynamics. *Progress in Aerospace Sciences*, 110:100542, October 2019.
- [38] Gaetan KW Kenway, Graeme J Kennedy, and Joaquim RRA Martins. Scalable parallel approach for high-fidelity steady-state aeroelastic analysis and adjoint derivative computations. *AIAA journal*, 52(5):935–951, 2014.
- [39] Anil Yildirim, Gaetan K. W. Kenway, Charles A. Mader, and Joaquim R. R. A. Martins. *Journal of Computational Physics*, page 108741, November.
- [40] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023.
- [41] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research*, 23(47):1–33, 2022.
- [42] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*, 2022.
- [43] Jascha Sohl-Dickstein, Eric A Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Advances in Neural Information Processing Systems*, pages 3516–3524, 2015.
- [44] Aravind Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *International Conference on Machine Learning*, pages 5843–5852, 2019.
- [45] Jonathan Ho, Xi Chen, Aditya Srinivas, Yan Duan, Pieter Abbeel, and Stefano Ermon. Denoising diffusion probabilistic models. In *International Conference on Learning Representations*, 2020.
- [46] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020.

- [47] Ming Chen Wang and George Eugene Uhlenbeck. On the theory of the brownian motion ii. *Reviews of modern physics*, 17(2-3):323, 1945.
- [48] H. Föllmer. An entropy approach to the time reversal of diffusion processes. In *Stochastic differential systems (Marseille-Luminy, 1984)*, volume 69 of *Lect. Notes Control Inf. Sci.*, pages 156–163. Springer, Berlin, 1985.
- [49] P. Cattiaux, G. Conforti, I. Gentil, and C. Léonard. Time reversal of diffusion processes under a finite entropy condition. *arXiv preprint arXiv:2209.11215*, Sept. 2022.
- [50] S. Chen, S. Chewi, J. Li, et al. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. *arXiv preprint arXiv:2209.11215*, 2022.
- [51] Alexander C Li, Mihir Prabhudesai, Shivam Duggal, Ellis Brown, and Deepak Pathak. Your diffusion model is secretly a zero-shot classifier. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2206–2217, 2023.
- [52] Cheng Lu, Huayu Chen, Jianfei Chen, Hang Su, Chongxuan Li, and Jun Zhu. Contrastive energy prediction for exact energy-guided diffusion sampling in offline reinforcement learning. In *International Conference on Machine Learning*, pages 22825–22855. PMLR, 2023.
- [53] K P Murphy. *Probabilistic machine learning: Advanced topics*. MIT press, 2023.
- [54] Huayu Chen, Cheng Lu, Chengyang Ying, Hang Su, and Jun Zhu. Offline reinforcement learning via high-fidelity generative behavior modeling. *arXiv preprint arXiv:2209.14548*, 2022.
- [55] Paul Adrien Maurice Dirac. *The principles of quantum mechanics*. Number 27. Oxford university press, 1981.
- [56] Wenhan Xia, Chengwei Qin, and Elad Hazan. Chain of lora: Efficient fine-tuning of language models via residual learning. *arXiv preprint arXiv:2401.04151*, 2024.
- [57] Arpita Gupta, Pratik Pawade, and Ramadoss Balakrishnan. Deep residual network and transfer learning-based person re-identification. *Intelligent Systems with Applications*, 16:200137, 2022.




Figure 9 – Wings and associated C_p contour of Baseline geometry, CFD-based and Data-driven optimized wings. Data-driven cases involve using 140,000, 10,000 LHS data and 500 Diffusion-based sampling Data when $\beta = 32$ and 64.