

A SINGLE-LOOP BAYESIAN OPTIMIZATION FRAMEWORK FOR MULTIDISCIPLINARY DESIGN OPTIMIZATION UNDER UNCERTAINTY

Bingran Wang[1], Dongjin Lee[2] & Gyeolhee Lee[2]

[1]School of Aerospace, Mechanical and Mechatronics Engineering, The University of Sydney
[2]Department of Automotive Engineering, Hanyang University

Abstract

Modern aerospace system design increasingly relies on high-fidelity, multidisciplinary simulation models to assess performance and reliability. In early-stage conceptual design, however, these models are often subject to significant uncertainty arising from operating conditions, environmental variability, and modelling assumptions. Accounting for such uncertainties within multidisciplinary design optimization (MDO) frameworks remains challenging, as conventional multidisciplinary design optimization under uncertainty (MDOUU) approaches typically rely on nested, double-loop structures that incur prohibitive computational cost. This limitation restricts the practical use of MDOUU in aerospace system design, particularly when every model evaluation is computationally expensive. This paper presents a single-loop Bayesian optimization framework for low- to moderate-dimensional MDOUU problems. The proposed approach constructs a probabilistic surrogate model in the joint design–uncertainty space and reformulates MDOUU as a sequential learning problem. Instead of performing repeated high-fidelity uncertainty propagation at each optimization iteration, the robust objective is estimated directly from the surrogate model. A Bayesian acquisition function is then used to select the next model evaluation based on its expected contribution to improving the estimated robust optimum. Consequently, uncertainty quantification and optimization are integrated within a unified learning framework, requiring only a single expensive model evaluation per iteration. The proposed framework is demonstrated on a scalable nonlinear analytical benchmark and an aero-structural wing design problem. Results show that the proposed method can consistently find more optimal design solutions than conventional surrogate-assisted MDOUU approaches under limited model evaluation budgets.

Keywords: multidisciplinary design optimization, optimization under uncertainty

1. Introduction

Aerospace system design is inherently multidisciplinary, requiring the simultaneous consideration of aerodynamics, structures, propulsion, flight dynamics, controls, and mission performance. Strong couplings between these disciplines mean that design decisions made in one area can have significant and often non-intuitive impacts on others—e.g., a small wing-area change propagates through the trim/weight coupling and shifts propulsion sizing. To address this complexity, multidisciplinary design optimization (MDO) has emerged as a powerful framework for systematically exploring design trade-offs while accounting for interdisciplinary interactions. Over the past decade, MDO has been successfully applied to a wide range of aerospace systems, including conventional transport aircraft [1, 2, 3], supersonic aircraft [4, 5], electric vertical take-off and landing (eVTOL) aircraft [6], spacecraft [7], satellites [8], and wind turbines [9, 10], enabling more efficient and higher-performance designs than traditional sequential single-discipline approaches that fix one discipline before optimizing.

Despite these successes, most existing MDO formulations remain deterministic, implicitly assuming that model inputs, operating conditions, and system parameters are known with certainty. In practice, aerospace systems are subject to significant sources of uncertainty arising from atmospheric variability, mission and operational conditions, manufacturing tolerances, material properties, and modeling assumptions. Ignoring these uncertainties during the design process can lead to designs that are overly conservative or, more critically, designs that fail to meet performance or safety requirements once deployed. Incorporating uncertainty directly into the optimization process leads to multidisciplinary design optimization under uncertainty (MDOUU), which enables designers to explicitly balance performance, robustness, and reliability.

The potential benefits of MDOUU for aerospace system design have been demonstrated in several recent studies [11, 12, 13, 14]. For example, in an eVTOL aircraft design problem [13], MDOUU was shown to reduce overall vehicle weight by more than 10% compared with safety-factor-based deterministic MDO, highlighting the substantial performance gains that can be achieved when design objectives explicitly account for uncertainty instead of masking it with ad-hoc safety factors. These results underscore the importance of MDOUU as a design paradigm for emerging aerospace systems, where tight performance margins and stringent certification requirements make robustness and reliability critical [15].

However, despite its conceptual appeal and demonstrated benefits, the practical adoption of MDOUU remains severely limited by the computational cost. Classical MDOUU formulations typically employ a double-loop structure, in which an outer optimization loop updates the design variables while an inner loop performs uncertainty quantification (UQ) to estimate a risk measure—mean, variance, quantile, or probability of failure. This nested structure leads to a significant growth in computational expense, making it prohibitively expensive when high-fidelity multidisciplinary models are used. Back in 2002, a NASA report in [16] identified the lack of efficient computational methods as a major barrier to the practical adoption of MDOUU in real-world aerospace system design problems. Unfortunately, this challenge remains unresolved for high-fidelity MDOUU.

To tackle this challenge, recent research has proposed several strategies to accelerate the UQ loop within MDOUU. These include adaptive sampling [17, 18],multi-fidelity methods [19], and computational graph transformation methods [20, 21]. These approaches can significantly reduce the cost of UQ evaluations in each optimization iteration; however, they continue to operate within the traditional double-loop MDOUU, where a full or partial UQ analysis must be performed at every optimization iteration. As a result, the inner UQ loop remains the dominant contributor to computational cost and continues to limit the scalability of MDOUU to high-fidelity aerospace design problems

Another related line of research seeks to reduce computational cost through surrogate-based optimization under uncertainty (OUU) [22, 23]. In these approaches, surrogate models are constructed to approximate the system response, and adaptive sampling strategies are employed to improve surrogate accuracy throughout the design–uncertainty space. While effective for reducing the cost of model evaluations, the adaptive sampling strategies all focus on improving the accuracy of the surrogate model rather than directly improving knowledge of the robust optimum. As a result, computational resources may be spent refining regions of the design space that have little influence on the final design decision.

These observations highlight a critical gap in the current literature: the absence of a practical framework that integrates uncertainty quantification and optimization into a single-loop process. Such a framework would avoid repeated inner-loop UQ analyses and instead make robust design decisions using stochastic information gathered progressively throughout the optimization process. Addressing this challenge is particularly important for early-stage aerospace design, where computational budgets are limited and rapid exploration of the design space is essential.

Motivated by this need, this paper presents a single-loop stochastic optimization framework for MDOUU. The proposed approach reformulates the MDOUU problem as a sequential learning task by constructing a probabilistic surrogate model over the combined design–uncertainty space. Rather than performing repeated high-fidelity uncertainty propagation at each optimization iteration, the robust objective is estimated directly from the surrogate model. At each iteration, the surrogate-based MDOUU problem is solved to identify the current robust optimum, and a Bayesian acquisition function

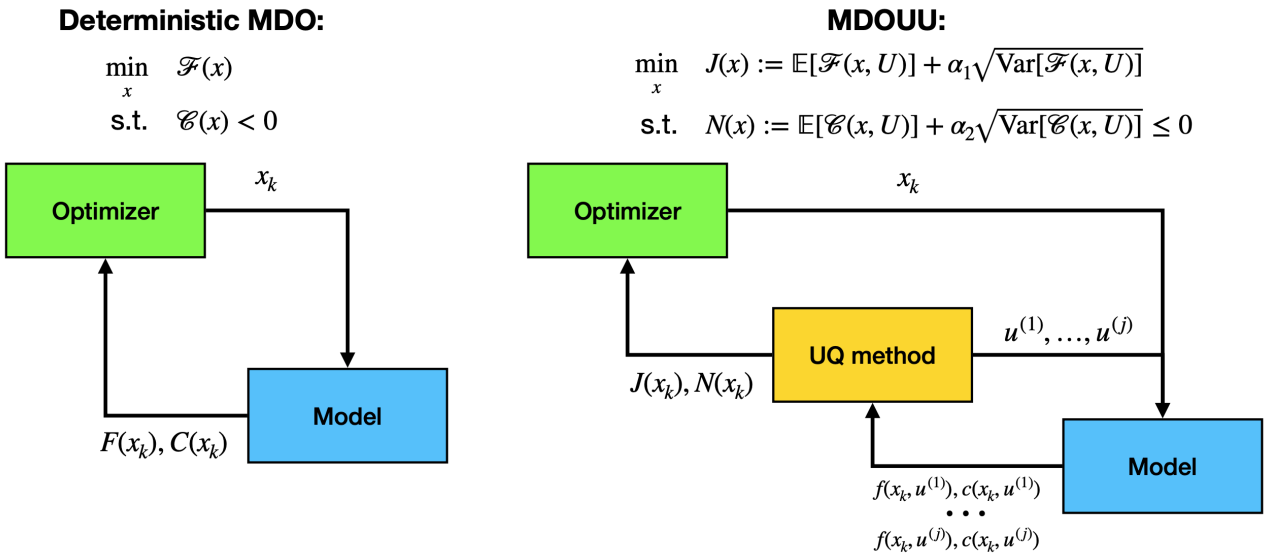


Figure 1 – Conventional computational frameworks for solving MDO and MDOUU problem.

is employed to select the next model evaluation in the joint design–uncertainty space. Unlike conventional surrogate-based OUU methods that primarily seek to improve global surrogate accuracy, the proposed acquisition strategy directly targets uncertainty in the robust optimum. By integrating UQ and optimization within a unified sequential learning framework, the proposed method eliminates the need for repeated inner-loop UQ analyses and substantially reduces the number of expensive model evaluations required to identify robust optimal designs.

2. Background

We begin by considering a deterministic *multidisciplinary design optimization (MDO)* problem formulated using the *multidisciplinary feasible (MDF)* architecture [24]:

$$\begin{aligned} \min_{x} \quad & \mathscr{F}(x) \\ \text{subject to} \quad & \mathscr{C}(x) \le 0, \end{aligned} \tag{1}$$

where $x \in \mathbb{R}^{n_x}$ denotes the vector of design variables, $\mathscr{F}(x)$ is the objective function, and $\mathscr{C}(x)$ represents a vector of inequality constraints. In the MDF architecture, disciplinary couplings are resolved internally at each function evaluation, such that the optimizer operates only on the design variables. The deterministic formulation in Eq. (1) can be extended to account for uncertainty in system inputs and parameters, leading to *multidisciplinary design optimization under uncertainty (MDOUU)*. Let Ω denote the sample space. We define $u \in \mathbb{R}^{n_u}$ as a realization of the random vector $U : \Omega \to \mathbb{R}^{n_u}$ with joint probability density $\rho(u)$. A commonly used robust design optimization (RDO) formulation expresses the objective and constraints in terms of statistical moments:

$$\begin{aligned} \min_{x} \quad & \mathbb{E}[\mathscr{F}(x,U)] + \alpha_1 \sqrt{\mathrm{Var}[\mathscr{F}(x,U)]} \\ \text{subject to} \quad & \mathbb{E}[\mathscr{C}(x,U)] + \alpha_2 \sqrt{\mathrm{Var}[\mathscr{C}(x,U)]} \le 0 \end{aligned} \tag{2}$$

where $\mathbb{E}[\cdot]$ and $\mathrm{Var}[\cdot]$ denote the expectation and variance operators, respectively. The coefficients $\alpha_1 \in \mathbb{R}_{\ge 0}$ and $\alpha_2 \in \mathbb{R}_{\ge 0}$ control the trade-off between nominal performance and robustness in the objective and constraints. Alternative formulations are also widely used, such as reliability-based design optimization (RBDO), in which the probability of constraint satisfaction is required to exceed a prescribed threshold.

In the remainder of this work, we focus on the robust objective in Eq. (2); the proposed framework extends naturally to robust constraints. Solving Eq. (2) typically requires a *double-loop computational structure*. At each outer optimization iteration, a forward *uncertainty quantification (UQ)* problem must be solved to evaluate the statistical quantities in the objective and constraints. In practice,

the inner loop is implemented using sampling- or quadrature-based methods, such as Monte Carlo simulation and polynomial chaos expansion [25, 26]. Even when surrogates are used inside the UQ loop, the nested structure still demands $\mathcal{O}(10^2 - 10^4)$ high-fidelity calls per outer iteration, which becomes prohibitive for expensive multidisciplinary analyses. The computational cost therefore grows rapidly with the dimension of the uncertainty space and can become prohibitive when expensive multidisciplinary analyses are required. A schematic comparison between computational frameworks for deterministic MDO and MDOUU is shown in Fig. 1.

3. Methodology

3.1 Problem formulation

We consider the MDOUU setting introduced in Eq. (2). In the numerical studies of this work, constraints are omitted and attention is restricted to the robust objective. In particular, we minimize

$$J(x) = \mathbb{E}[\mathscr{F}(x,U)] + \alpha\sqrt{\mathrm{Var}[\mathscr{F}(x,U)]}, \tag{3}$$

where $x \in \mathbb{R}^{n_x}$ denotes the design variables, $U : \Omega \to \mathbb{R}^{n_u}$ denotes the uncertain inputs, and $\alpha \geq 0$ is a user-defined robustness parameter.

In a conventional double-loop implementation, evaluating $J(x)$ at a candidate design requires repeated high-fidelity evaluations of $\mathscr{F}(x,u)$ over u to propagate the uncertain inputs. The proposed framework avoids this repeated inner-loop statistical moment estimation on the expensive model by learning $\mathscr{F}(x,u)$ from a growing dataset of joint realizations and performing robust optimization on a probabilistic surrogate instead.

3.2 Single-loop Bayesian optimization framework

The key modeling choice is to construct a global surrogate over the combined design–uncertainty space

$$z := (x,u) \in \mathscr{Z} := \mathscr{X} \times \mathscr{U} \subset \mathbb{R}^{n_x+n_u}, \tag{4}$$

and to treat MDOUU as a sequential learning problem on $\mathscr{Z}$. Rather than propagating uncertainty on the high-fidelity model at fixed x, the algorithm collects N_k paired observations

$$y_i = \mathscr{F}(x_i,u_i), \qquad i = 1,\ldots,N_k, \tag{5}$$

and updates a probabilistic surrogate after each new sample. The response surface

$$(x,u) \mapsto \mathscr{F}(x,u) \tag{6}$$

is therefore learned directly in joint space, while the robust objective and its minimizer are estimated from the surrogate at every iteration. UQ and design optimization are integrated within a single loop, and each iteration requires only one model evaluation.

Figure 2 summarizes the workflow; Algorithm 1 gives a formal description. After initialization, the procedure maintains a dataset

$$\mathscr{D}_k = \{(x_i,u_i,y_i)\}_{i=1}^{N_k}. \tag{7}$$

Each iteration performs the following steps until a stopping criterion is met or an evaluation budget is exhausted. We write $\hat{\mathscr{F}}_k$ for the surrogate fitted at iteration k and $\tilde{J}_k(x)$ for the robust objective in Eq. (9) evaluated under $\hat{\mathscr{F}}_k$.

Step 1: Surrogate update. Fit a probabilistic surrogate $\hat{\mathscr{F}}_k$ to $\mathscr{D}_k$.

Step 2: Surrogate robust optimization. Compute the surrogate-based robust minimizer x_k^* and value $v_k = \tilde{J}_k(x_k^*)$.

Step 3: Acquisition. Select the next joint point $(x_{k+1},u_{k+1}) \in \mathscr{Z}$.

Step 4: High-fidelity evaluation. Evaluate $y_{k+1} = \mathscr{F}(x_{k+1},u_{k+1})$ and augment $\mathscr{D}_k$.

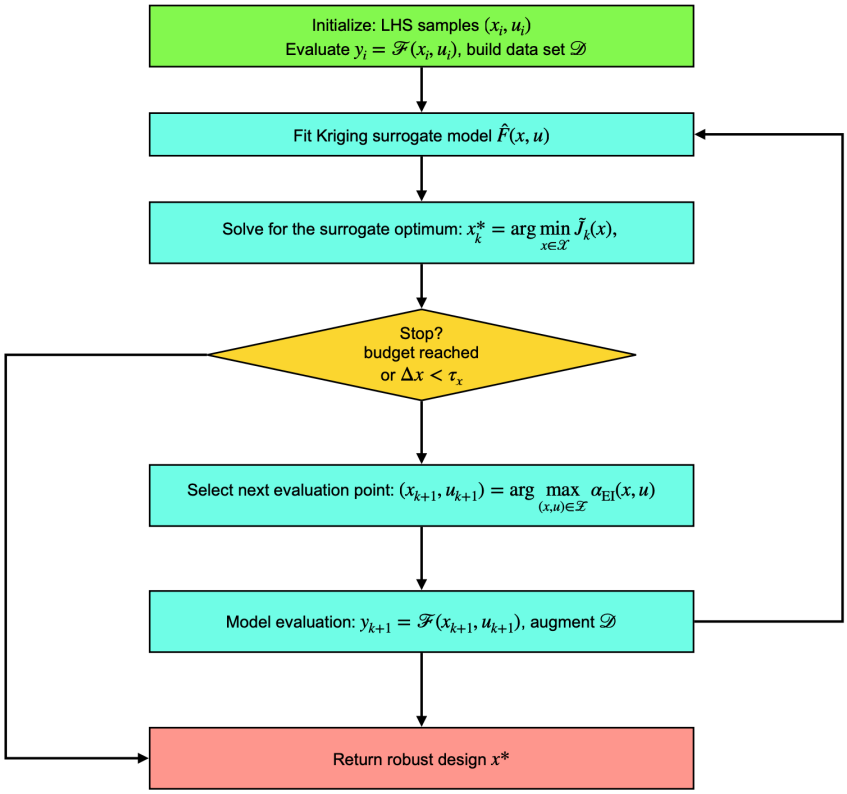


Figure 2 – Overview of the proposed single-loop Bayesian optimization framework for MDOUU.

Initialization draws N_0 space-filling samples $\{(x_i, u_i)\}_{i=1}^{N_0}$ by Latin hypercube sampling (LHS) on $\mathscr{X}$ and evaluates the expensive model at each point. In the present study, N_0 is a small fraction of the total evaluation budget.

Unlike double-loop MDOUU, the true objective $J(x)$ in Eq. (3) is not recomputed on the high-fidelity model at every iteration. Instead, the algorithm tracks the evolution of the surrogate robust minimizer x_k^* and allocates new evaluations in $\mathscr{X}$ where the gradient-based expected improvement acquisition (introduced in Sect. 3.4) predicts the largest expected reduction in $|x_k^* - x_{k-1}^*|$.

3.3 Probabilistic surrogate and surrogate robust objective

Given $\mathscr{D}_k$, a Kriging (Gaussian process) surrogate is fitted on $z = (x, u)$. For any fixed input (x, u), the posterior predictive distribution is Gaussian,

$$\mathscr{F}(x,u) \mid \mathscr{D}_k \sim \mathscr{N}\left(\mu_f(x,u), \sigma_f^2(x,u)\right), \tag{8}$$

where μ_f and σ_f denote the predictive mean and standard deviation. Kriging is adopted because it provides a differentiable mean predictor $\mu_f(x,u)$ together with an explicit epistemic uncertainty estimate $\sigma_f(x,u)$, both of which are required for gradient-based acquisition in joint space (presented in Sect. 3.4).

It is important to distinguish Eq. (8) from the distribution of the random variable $\mathscr{F}(x,U)$ when U is aleatoric. Equation (8) is a *pointwise* predictive distribution. For fixed x, the random quantity $\mathscr{F}(x,U)$ induced by U is generally *non-Gaussian*, because the maps $u \mapsto \mu_f(x,\cdot)$ and $u \mapsto \sigma_f(x,\cdot)$ are typically nonlinear, and including epistemic uncertainty yields a mixture over U. Hence, the true robust objective $J(x)$ in Eq. (3) has no closed-form expression in terms of μ_f and σ_f alone.

We therefore approximate $J(x)$ on the surrogate by propagating only the predictive moments through the mean–standard-deviation structure of Eq. (3):

$$\tilde{J}_k(x) = \mathbb{E}_U[\mu_f(x,U)] + \beta \sqrt{\mathrm{Var}_U[\mu_f(x,U)] + \beta_e^2\, \mathbb{E}_U\left[\sigma_f^2(x,U)\right]}, \tag{9}$$

where β is the robustness weight on the surrogate objective (typically $\beta = \alpha$), and $\beta_e \geq 0$ controls the contribution of epistemic predictive variance. The first term is the mean of the surrogate mean

predictor over U; the second term under the square root combines aleatoric variability of $\mu_f(x,U)$ with a weighted average of local epistemic variance $\sigma_f^2(x,U)$. Setting $\beta_e = 0$, as in the numerical experiments below, gives

$$\tilde{J}_k(x) = \mathbb{E}_U[\mu_f(x,U)] + \beta\sqrt{\mathrm{Var}_U[\mu_f(x,U)]}, \tag{10}$$

i.e., a mean–standard-deviation objective applied to $\mu_f(x,U)$ only. Setting $\beta_e = 1$ recovers a law-of-total-variance approximation, $\mathrm{Var}_U[\mathscr{F}(x,U)] \approx \mathrm{Var}_U[\mu_f(x,U)] + \mathbb{E}_U[\sigma_f^2(x,U)]$, before taking the square root.

The operators $\mathbb{E}_U[\cdot]$ and $\mathrm{Var}_U[\cdot]$ are taken with respect to the distribution of U only. We evaluate them with N_u paired nodes and weights $\{(u_j, w_j)\}_{j=1}^{N_u}$ matched to $\rho(u)$ — Gauss–Legendre quadrature for uniform U, Gauss–Hermite for Gaussian U, or Monte Carlo samples $u_j \sim \rho(u)$ with $w_j = 1/N_u$:

$$\mathbb{E}_U[\mu_f(x,U)] = \sum_{j=1}^{N_u} w_j\,\mu_f(x,u_j), \tag{11}$$

$$\mathrm{Var}_U[\mu_f(x,U)] = \sum_{j=1}^{N_u} w_j\,\mu_f(x,u_j)^2 - \left(\sum_{j=1}^{N_u} w_j\,\mu_f(x,u_j)\right)^2, \tag{12}$$

$$\mathbb{E}_U\left[\sigma_f^2(x,U)\right] = \sum_{j=1}^{N_u} w_j\,\sigma_f^2(x,u_j). \tag{13}$$

Because surrogate evaluations are inexpensive, these sums are computed at negligible cost relative to one expensive high-fidelity evaluation of $\mathscr{F}$.

At iteration k, the surrogate-based robust optimum is defined as

$$x_k^* = \arg\min_{x\in\mathscr{X}} \tilde{J}_k(x), \qquad v_k := \tilde{J}_k(x_k^*). \tag{14}$$

Equation (14) is solved by gradient-based optimization methods with multiple random restarts in x, warm-started from x_{k-1}^* when available. The pair (x_k^*, v_k) represents the algorithm's current belief about the robust optimum and provides the reference state for acquisition and stopping.

3.4 Acquisition in joint design–uncertainty space

The surrogate is trained on joint samples (x_i, u_i), meaning the next evaluation must be selected within the joint space $\mathscr{Z}$ rather than the design space $\mathscr{X}$ alone. The acquisition function should therefore identify realizations that are informative for refining the estimate of the robust optimum.

Robust-objective acquisition. The acquisition value at a candidate (x,u) can be defined as the expected improvement in the surrogate-based robust objective after incorporating a new observation:

$$\alpha_{\mathrm{rob}}(x,u) = \mathbb{E}\left[\max\left(v_k - v_{k+1}, 0\right)|(x,u)\right], \tag{15}$$

where v_k is the current optimal robust objective value, and v_{k+1} is the updated value obtained by refitting the surrogate after a probabilistic observation at (x,u) and recomputing Eq. (14). Although Equation (15) directly targets improvement in the robust optimum via predicted observations, its evaluation is computationally intensive, requiring repeated surrogate refitting and nested inner robust optimizations over a candidate set.

Gradient-based expected improvement. We employ an efficient, gradient-based expected improvement (EI) acquisition function on the surrogate response:

$$\alpha_{\mathrm{EI}}(x,u) = \mathbb{E}\left[\max\left(f_{\min} - Y, 0\right)\right], \qquad Y \sim \mathscr{N}\left(\mu_f(x,u), \sigma_f^2(x,u)\right), \tag{16}$$

where $f_{\min} = \min_i y_i$ is the best observed value in $\mathscr{D}_k$. Under the Gaussian surrogate assumption, α_{EI} admits a closed-form expression. The next sample is obtained from

Algorithm 1 Bayesian optimization under uncertainty

Require: $\mathscr{X}$, $\mathscr{U}$; weights α, β; initial size N_0; post-init minimum $N_{\min}$; tolerance τ_x; budget $N_{\max}$

Sample $\{(x_i, u_i)\}_{i=1}^{N_0}$ by LHS on $\mathscr{X} \times \mathscr{U}$
Evaluate $y_i \leftarrow \mathscr{F}(x_i, u_i)$ and set $\mathscr{D} \leftarrow \{(x_i, u_i, y_i)\}_{i=1}^{N_0}$
Fit surrogate $\hat{\mathscr{F}}$ to $\mathscr{D}$
$x^* \leftarrow \arg\min_{x \in \mathscr{X}} \tilde{J}(x)$ ▷ Eq. (9)
while $|\mathscr{D}| < N_{\max}$ **do**
    Compute Δx from successive minimizers via Eq. (18)
    if $\Delta x < \tau_x$ and $|\mathscr{D}| \geq N_0 + N_{\min}$ **then**
        break
    end if
    $(x, u) \leftarrow \arg\max_{(x,u) \in \mathscr{Z}} \alpha_{\mathrm{EI}}(x, u)$ ▷ warm-start at (x^*, u_{init})
    $y \leftarrow \mathscr{F}(x, u)$; augment $\mathscr{D}$
    Refit $\hat{\mathscr{F}}$; update $x^* \leftarrow \arg\min_{x \in \mathscr{X}} \tilde{J}(x)$
end while
return x^*

$$(x_{k+1}, u_{k+1}) = \arg \max_{(x,u) \in \mathscr{Z}} \alpha_{\mathrm{EI}}(x, u), \tag{17}$$

using the Sequential Least Squares Programming (SLSQP) algorithm on $z = (x, u)$ with N_{rs} multiple warm starts ($N_{\mathrm{rs}} = 10$ in our experiments). The design variable component is initialized at x_k^*, while the random variable component is initialized from a candidate pool that includes $u = \mu_u$, the random sample maximizing $\sigma_f(x_k^*, u)$, and additional random draws from $\mathscr{U}$.
Although α_{EI} targets pointwise improvement of the response $\mathscr{F}$ rather than of the robust objective $\tilde{J}$, coupling the warm start at x_k^* with the inner solve in Eq. (14) yields a practical single-loop strategy: each expensive evaluation refines the surrogate in joint space (x, u) near x_k^*, while acquisition is initialized at the current surrogate robust minimizer.

3.5 Stopping criterion

As the surrogate improves, the predicted robust minimizer x_k^* is expected to stabilize. Convergence is therefore monitored through the relative change

$$\Delta x_k = \frac{\left\| x_k^* - x_{k-1}^* \right\|_2}{\left\| x_{k-1}^* \right\|_2 + \varepsilon}, \tag{18}$$

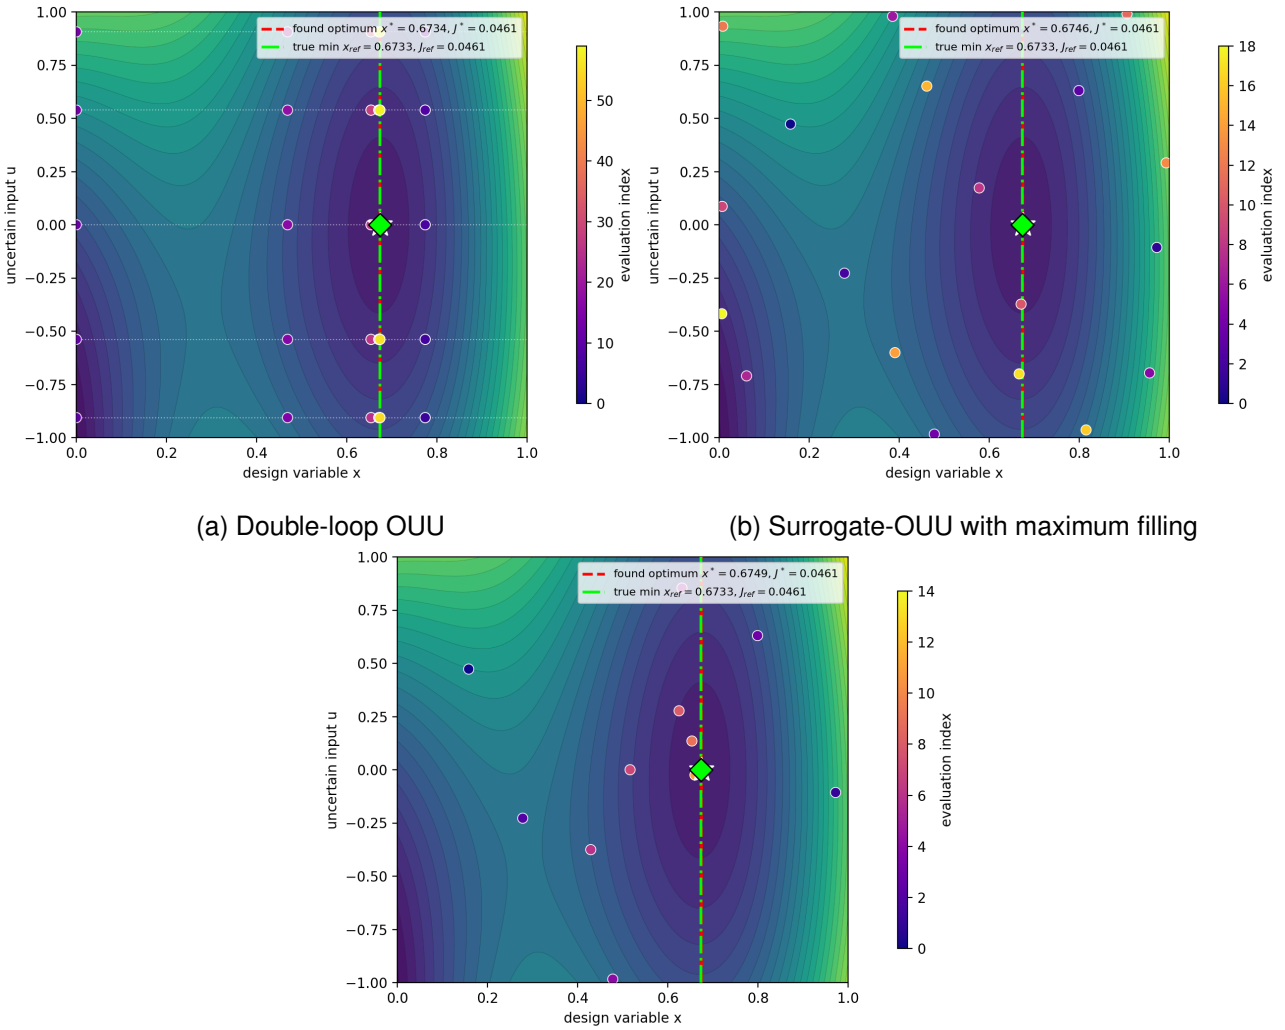
where $\varepsilon > 0$ prevents division by zero. The loop terminates when

$$\Delta x_k < \tau_x \tag{19}$$

for a user-defined tolerance τ_x, provided that a minimum number of post-initialization evaluations has been completed; otherwise, termination occurs when the evaluation budget $N_{\max}$ is reached. The final design is taken as x_k^* from the last iteration.

4. Numerical results

We refer to the proposed method as Bayesian Optimization Under Uncertainty (BOUU) and evaluate its performance on (a) a scalable analytical benchmark with a reference robust optimum (Sect. 4.1) and (b) the NASA wing-aerostructural problem with cruise-condition uncertainty (Sect. 4.2). The analytical case isolates evaluation-efficiency vs. a double-loop reference; the wing case tests scalability when high-fidelity evaluations dominate cost.



(a) Double-loop OUU

(b) Surrogate-OUU with maximum filling

(c) Proposed BOUU

Figure 3 – 1D demonstration: contour of $\mathscr{F}(x,u)$ with evaluation history (color indicates evaluation index), true robust minimum (green), and found optimum (red).

4.1 Analytical problems

We first demonstrate the proposed framework on a scalable nonlinear analytical benchmark with uniform aleatoric uncertainty. Figure 3 previews the one-dimensional case; the problem definition, reference solution, and quantitative comparison follow below. The analytical function involved is

$$\mathscr{F}(x,u) = \sum_{i=1}^{n_x} \left[(x_i - 0.3)^2 + 0.2\sin(6x_i)\right] + \frac{0.5}{n_u}\sum_{j=1}^{n_u} u_j\,(x-0.7)^2 + \frac{0.1}{n_u}\sum_{j=1}^{n_u} u_j^2, \tag{20}$$

where $x \in [0,1]^{n_x}$ are design variables, $u \in [-1,1]^{n_u}$ are uncertain inputs. Eq. (20) reduces to the one-dimensional test function used in the contour demonstrations. The OUU can be described as

$$\min_{x} \quad J(x) := \mathbb{E}[\mathscr{F}(x,U)] + \sqrt{\mathrm{Var}[\mathscr{F}(x,U)]} \tag{21}$$

with U distributed uniformly on $[-1,1]^{n_u}$.
A reference solution $(x_{\mathrm{ref}}, J_{\mathrm{ref}})$ is computed offline for each dimension using a dense search in x and high-accuracy uncertainty propagation in u. All reported errors are measured against this reference. Three methods are compared:

(a) **Double-loop OUU:** tensor-product five-point Gauss–Legendre inner UQ (5^{n_u} quadrature nodes) on the true model, with SLSQP and finite-difference gradients in the outer loop;

Table 1 – One-dimensional analytical benchmark result ($n_x = n_u = 1$).

Method	x^*	J^*	$\lvert J^* - J_{\text{ref}}\rvert$	Evaluations
Reference	0.6733	0.04606	—	—
Double-loop OUU	0.6734	0.04606	8.3×10^{-8}	60
Surrogate-OUU with LHS filling	0.6746	0.04606	4.8×10^{-6}	19
Proposed BOUU	0.6749	0.04607	7.6×10^{-6}	**15**

(b) **Surrogate-OUU with maximum space-filling:** surrogate-assisted OUU with maximin joint sampling;

(c) **Proposed BOUU:** the single-loop MDOUU method proposed in this work.

4.1.1 One-dimensional demonstration

We first compare the three methods on the $n_x = n_u = 1$ case. Figure 3 shows the joint response $\mathscr{F}(x,u)$, the sequence of evaluated points, the true robust minimum, and the design identified by each method. Table 1 summarizes the 1D results.

Overall, all three methods identify the robust optimum accurately. The proposed BOUU attains comparable objective accuracy with the fewest expensive evaluations (15 vs. 60 for double-loop OUU), reflecting the efficiency of informative sampling in joint (x,u) space. The double-loop approach requires more evaluations because each outer step performs repeated inner-loop uncertainty propagation on the high-fidelity model.

Comparing the sampling trajectories in Fig. 3 highlights the structural difference among the methods. Double-loop OUU concentrates evaluations along vertical lines in (x,u) at the current design iterate, consistent with fixed quadrature nodes in u at each outer update of x. Surrogate-OUU with maximum space-filling spreads points across the joint domain to build a global surrogate before refining the robust design estimate. In contrast, BOUU places fewer evaluations and progressively focuses on the neighborhood of the emerging robust optimum, where acquisition and surrogate-based inner optimization interact within a single loop.

4.1.2 Scalability under fixed evaluation budget

To assess scalability, Methods (b) and (c) are compared on the scalable problem with $n_x = n_u \in \{2,3,4,5\}$ under a *fixed* expensive-evaluation budget of $N_{\max} = 50$. Early stopping is disabled ($\tau_x = 0$) so that each run uses the full budget. Surrogate initialization is fixed at $N_0 = 10$ for all dimensions. For each (d, method) pair, 20 independent random seeds are used; Figure 4 reports the median absolute objective error

$$\lvert J^* - J_{\text{ref}}\rvert$$

with error bars indicating the interquartile range. Figure 4 shows that the proposed BOUU achieves lower median objective error than Surrogate-OUU with maximum filling at every tested dimension.

These results support the central claim of the proposed framework: under a strict and equal evaluation budget, targeting informative joint samples with acquisition guided by the evolving surrogate robust minimizer yields consistently better robust designs than passive space-filling on the surrogate alone as dimension increases from two to five.

4.2 Aero-structural wing design optimization

The second benchmark is a NASA Common Research Model (CRM) wing design MDOUU problem built in OpenAeroStruct [27]. The model couples a vortex-lattice aerodynamic analysis to a tubular-spar structural finite-element model with two-way load and displacement transfer resolved at each flight condition. The design goal is to identify wing twist and spar-thickness distributions that minimize penalized cruise fuel burn while remaining structurally feasible over uncertain cruise altitude and Mach number. Each evaluation requires a nonlinear aerostructural solution with inner trim for $L = W$, and the response can be strongly nonlinear when structural failure penalties are activated.

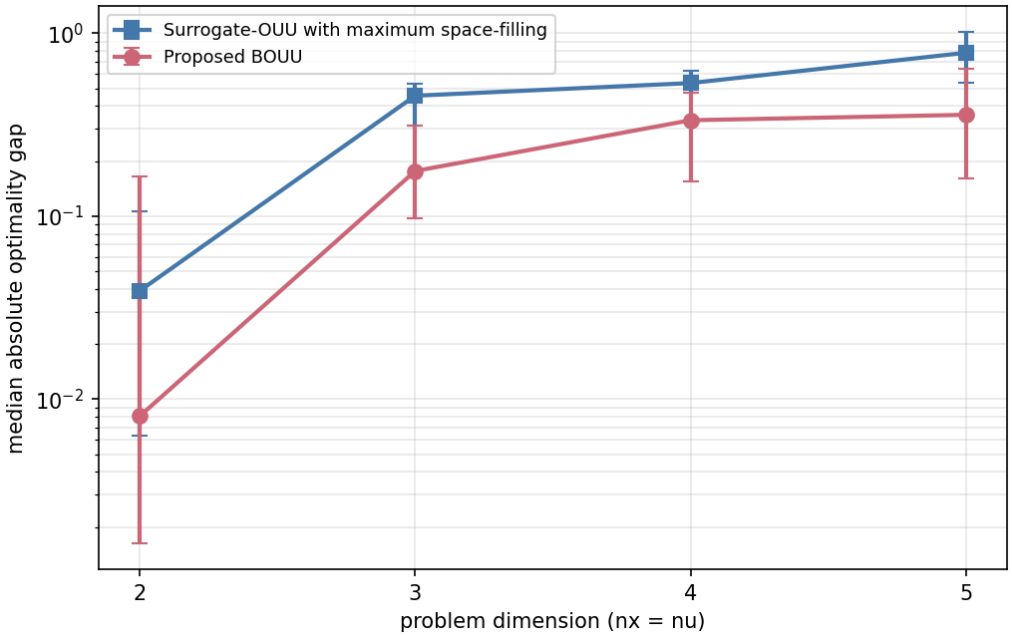


Figure 4 – Scalability result on the analytical benchmark problem: median absolute optimality gap ($|J^* - J_{\text{ref}}|$) versus problem dimension ($n_x = n_u$). Fixed evaluation budget $N_{\max} = 50$, initial design size $N_0 = 10$, and 20 random seeds per method and dimension; error bars show the interquartile range.

Table 2 – Aero-structural wing MDOUU formulation.

	Description	Range (units)	Quantity
Objective	Fuel burn + penalized structural constraints	$J(x) = \mu_{\mathscr{F}} + \sigma_{\mathscr{F}}$ $\mathscr{F} = m_{\text{fuel}} + \rho_{\text{fail}} \max(0, g_{\text{fail}})^2$	1
DESIGN VARIABLES	Wing twist control points	$\theta_{\text{cp}} \in [-10, 15]$ (deg)	5
	Spar thickness control points	$\mathbf{t}_{\text{cp}} \in [0.01, 0.5]$ (m)	3
RANDOM VARIABLES	Cruise altitude	$\mathscr{U}(31{,}000, 39{,}000)$ (ft)	1
	Cruise Mach number	$\mathscr{U}(0.80, 0.88)$	1

The MDOUU formulation is summarized in Table 2. We formulated a low-dimensional MDOUU problem, with eight design variables and two aleatoric uncertain inputs to test the performance of our proposed method in low-dimensional settings. The design vector contains five B-spline twist control points and three spar thickness control points. The variations in operating conditions are characterized as uniform random variables mapped to cruise altitude in $[31{,}000, 39{,}000]$ ft and Mach number in $[0.80, 0.88]$. At each sampled (x, u) pair, angle of attack is trimmed internally so that lift balances aircraft weight at the realized atmospheric state. The quantity of interest is a penalized fuel-burn metric that combines cruise performance with structural feasibility. The primary term, $m_{\text{fuel}}(x, u)$, is the cruise fuel burn (kg) returned by the OpenAeroStruct performance model after the coupled aerostructural analysis and trim solve; it depends on both the aerodynamic loads at the sampled flight condition and the wing structural weight associated with design x. The secondary term penalizes violation of the wing structural failure criterion. OpenAeroStruct reports an aggregated failure quantity $g_{\text{fail}}(x, u)$ obtained from the tubular-spar von Mises stresses using a Kreisselmeier–Steinhauser (KS) function; a design is structurally feasible at the sampled condition when

$$g_{\text{fail}}(x, u) \leq 0. \tag{22}$$

In the deterministic OpenAeroStruct formulation, this quantity is enforced as an explicit constraint. In the present MDOUU benchmark, however, the expensive multidisciplinary model is treated as a black-box evaluator inside the surrogate-based optimization loop, and structural feasibility is instead

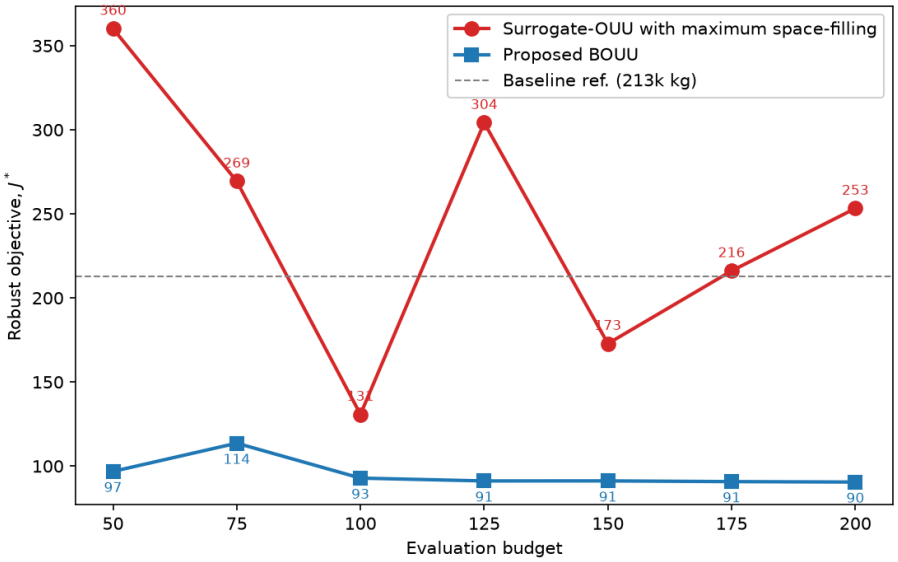


Figure 5 – Best robust objective J^* versus evaluation budget for Surrogate-OUU and proposed BOUU. Dashed line: baseline CRM design.

enforced through a quadratic penalty appended to the fuel-burn objective:

$$\mathscr{F}(x,u) = m_{\text{fuel}}(x,u) + \rho_{\text{fail}} \max\left(0, g_{\text{fail}}(x,u)\right)^2, \tag{23}$$

with penalty weight $\rho_{\text{fail}} = 10^5$. When $g_{\text{fail}} \leq 0$, the penalty term vanishes and $\mathscr{F} = m_{\text{fuel}}$. When the structure is overstressed, the penalty grows rapidly with constraint violation and dominates the response, discouraging the optimizer from exploring infeasible regions of the design space while preserving a scalar QoI suitable for surrogate modeling and robust objective estimation. This penalized formulation is particularly important in the CRM benchmark because aggressive thickness reductions can produce large fuel-burn improvements at nominally infeasible designs, and the failure penalty prevents the OUU methods from being misled by artificially low fuel values. The robust objective follows Eq. (3) with $\alpha = 1$ and uniform uncertainty on u. A coarse CRM discretization ($N_x = 2$, $N_y = 5$) keeps the benchmark affordable while retaining the essential aerostructural coupling. The baseline CRM geometry yields $J = 2.13 \times 10^5$ kg under this robust metric.

Methods (b) (Surrogate-OUU with maximin space-filling) and (c) (proposed BOUU) are compared under fixed evaluation budgets $N_{\max} \in \{50, 75, 100, 125, 150, 175, 200\}$. Both methods use the same Kriging-with-PLS surrogate method [28] from the SMT surrogate modeling toolbox [29], early stopping is disabled ($\tau_x = 0$) so that each run consumes its full budget. Figure 5 reports the best robust objective J^* attained at the end of each run, together with the baseline CRM reference. The proposed BOUU consistently finds designs with lower J^* than the baseline CRM geometry across all tested budgets. At $N_{\max} = 50$, BOUU reduces J^* from 2.13×10^5 kg to 9.68×10^4 kg, whereas Surrogate-OUU remains far above the baseline at 3.60×10^5 kg. As the budget increases, BOUU improves from 9.28×10^4 kg at $N_{\max} = 100$ to 9.05×10^4 kg at $N_{\max} = 200$, indicating stable convergence toward a robust wing design despite the limited evaluation budget. Surrogate-OUU does not exhibit comparable improvement: its best J^* varies between 1.31×10^5 and 3.60×10^5 kg over the tested budgets and remains above the baseline in several cases. This behavior matches the analytical study, suggesting that passive joint space-filling is insufficient when the feasible region is narrow and failure penalties distort the response surface in eight design dimensions.

Figure 6 visualizes the best BOUU design from the $N_{\max} = 200$ run evaluated at the mean flight condition ($h = 35{,}000$ ft, $M = 0.84$). Relative to the baseline geometry, the optimized design redistributes twist and spar thickness along the semispan and produces a deformed planform with lower penalized fuel burn at the nominal atmosphere. The structural response remains below the allowable von Mises stress at this condition, illustrating that the robust optimizer is balancing fuel performance with structural feasibility under flight-condition variability. Taken together, these results extend the analytical

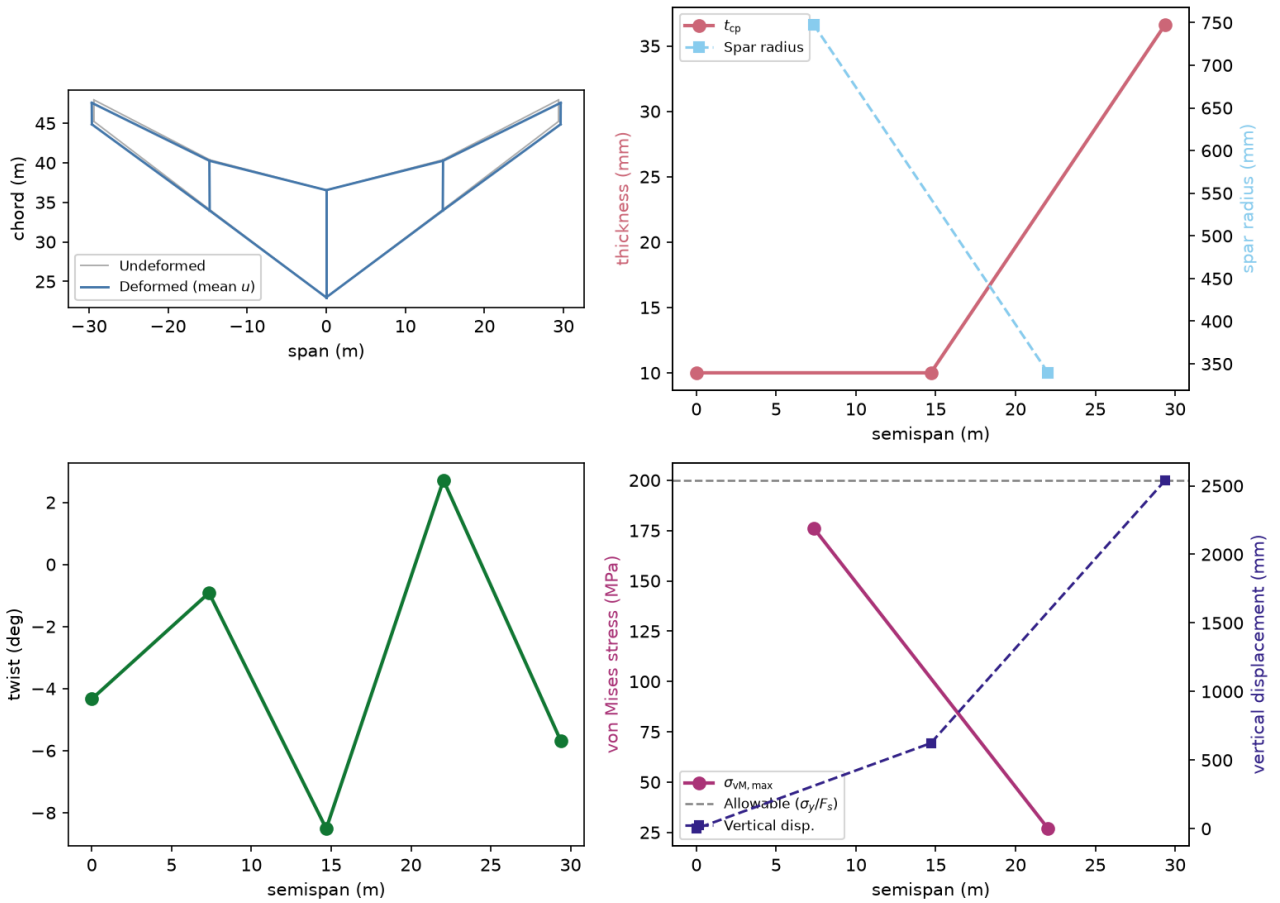


Figure 6 – Best BOUU design ($N_{\max} = 200$) at the mean flight condition.

benchmark findings to a realistic wing design problem and demonstrate that acquisition guided by the surrogate robust objective remains effective.

5. Conclusion

This paper addressed a central computational bottleneck in multidisciplinary design optimization under uncertainty (MDOUU): the repeated uncertainty quantification (UQ) required at every optimization iteration in the conventional double-loop structure. We presented a single-loop Bayesian optimization framework that learns a probabilistic surrogate of the model response in the joint space of design variables and uncertain inputs, estimates the robust objective on that surrogate, and selects each new expensive evaluation using an acquisition criterion tied to the evolving surrogate robust minimizer. By integrating uncertainty quantification and optimization within one learning loop, the proposed approach requires at most one high-fidelity multidisciplinary evaluation per optimization iteration and eliminates the need for an inner UQ loop on the expensive model during optimization.

The framework was evaluated on a scalable nonlinear analytical benchmark and on a NASA Common Research Model wing aerostructural problem formulated in OpenAeroStruct with uncertain cruise altitude and Mach number. On the analytical problems, the proposed Bayesian optimization under uncertainty (BOUU) method achieved lower median absolute optimality gaps than a surrogate-based space-filling alternative under a fixed evaluation budget, with consistent advantages as problem dimension increased from two to five. On the aerostructural benchmark, BOUU consistently outperformed the space-filling strategy across evaluation budgets from fifty to two hundred and identified wing designs with substantially lower robust penalized fuel burn than the baseline CRM geometry, whereas passive joint sampling remained unstable in the higher-dimensional, penalty-dominated response landscape. Together, these results demonstrate that informative acquisition in joint (x, u) space is an effective strategy for budget-limited MDOUU when multidisciplinary analyses dominate computational cost.

While the present study focused on low- to moderate-dimensional problems with a scalar robust objective and soft constraint handling through penalty terms, future work will scale to larger design and random spaces and explicit robust constraints, which will require further study on surrogate modeling and acquisition designs.

6. Contact Author Email Address

Corresponding author: Bingran Wang bingran.wang@sydney.edu.au

7. 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] Gaetan KW Kenway and Joaquim RRA Martins. Multipoint high-fidelity aerostructural optimization of a transport aircraft configuration. *Journal of Aircraft*, 51(1):144–160, 2014.

[2] L Leifsson, Andy Ko, William H Mason, Joseph A Schetz, Bernard Grossman, and Raphael T Haftka. Multidisciplinary design optimization of blended-wing-body transport aircraft with distributed propulsion. *Aerospace Science and Technology*, 25(1):16–28, 2013.

[3] Yiyuan Ma, Wei Zhang, and Ali Elham. Multidisciplinary design optimization of twin-fuselage aircraft with boundary-layer-ingesting distributed propulsion. *Journal of Aircraft*, 59(6):1588–1602, 2022.

[4] Sabet Seraj and Joaquim RRA Martins. Minimum trim drag for a three-surface supersonic transport aircraft. *Journal of Aircraft*, 62(3):719–730, 2025.

[5] Seongim Choi, Juan J Alonso, and Ilan M Kroo. Two-level multifidelity design optimization studies for supersonic jets. *Journal of Aircraft*, 46(3):776–790, 2009.

[6] Marius L Ruh, Michael AP Warner, Luca Scotzniovsky, Andrew H Fletcher, Mark Z Sperry, and John T Hwang. System-level, large-scale multidisciplinary design optimization of an air taxi concept. *Journal of Aircraft*, pages 1–24, 2026.

[7] Mathieu Balesdent, Nicolas Bérend, Philippe Dépincé, and Abdelhamid Chriette. A survey of multidisciplinary design optimization methods in launch vehicle design. *Structural and Multidisciplinary optimization*, 45(5):619–642, 2012.

[8] Masafumi Isaji, Yuji Takubo, and Koki Ho. Multidisciplinary design optimization approach to integrated space mission planning and spacecraft design. *Journal of Spacecraft and Rockets*, 59(5):1660–1670, 2022.

[9] Andrew Ning and K Dykes. Understanding the benefits and limitations of increasing maximum rotor tip speed for utility-scale wind turbines. In *Journal of physics: conference series*, volume 524, page 012087. IOP Publishing, 2014.

[10] Marco Mangano, Sicheng He, Yingqian Liao, Denis-Gabriel Caprace, Andrew Ning, and Joaquim RRA Martins. Wind turbine rotor design using high-fidelity aerostructural optimization. *AIAA Journal*, 63(9):3493–3513, 2025.

[11] Laure Jaeger, Christian Gogu, Stéphane Segonds, and Christian Bes. Aircraft multidisciplinary design optimization under both model and design variables uncertainty. *Journal of Aircraft*, 50(2):528–538, 2013.

[12] Daejin Lim and Kwanjung Yee. Sizing for advanced air mobility under uncertainty environment. In *AIAA AVIATION 2023 Forum*, page 4265, 2023.

[13] Bingran Wang, Marius L Ruh, Aoran Tian, Luca Scotzniovsky, and John T Hwang. Large-scale mdo under uncertainty of an evtol aircraft using dimension reduction via global sensitivity analysis. In *AIAA AVIATION FORUM AND ASCEND 2025*, page 3344, 2025.

[14] Bingran Wang, Nicholas C Orndorff, Anugrah J Joshy, and John T Hwang. Graph-accelerated large-scale multidisciplinary design optimization under uncertainty of a laser-beam-powered aircraft. In *AIAA SCITECH 2024 Forum*, page 0169, 2024.

[15] Debiao Meng, Shiyuan Yang, Chao He, Hongtao Wang, Zhiyuan Lv, Yipeng Guo, and Peng Nie. Multidisciplinary design optimization of engineering systems under uncertainty: a review. *International Journal of Structural Integrity*, 13(4):565–593, 2022.

[16] Thomas A Zang. *Needs and opportunities for uncertainty-based multidisciplinary design methods for aerospace vehicles*. National Aeronautics and Space Administration, Langley Research Center, 2002.

[17] Leo WT Ng and Karen E Willcox. Monte carlo information-reuse approach to aircraft conceptual design optimization under uncertainty. *Journal of Aircraft*, 53(2):427–438, 2016.

[18] Garo Bedonian, Jason E Hicken, and Edwin Forster. Robust design optimization with gradient-based surrogates and adaptive sampling: g. bedonian et al. *Structural and Multidisciplinary Optimization*, 68(8):166, 2025.

[19] Jaeyub Hyun, Anirban Chaudhuri, Karen E Willcox, and H Alicia Kim. Robust topology optimization of continuum structures under material uncertainties using multifidelity monte carlo. *Computer Methods in Applied Mechanics and Engineering*, 451:118714, 2026.

[20] Bingran Wang, Mark Sperry, Victor E Gandarillas, and John T Hwang. Accelerating model evaluations in uncertainty propagation on tensor grids using computational graph transformations. *Aerospace Science and Technology*, 145:108843, 2024.

[21] Bingran Wang, Nicholas C Orndorff, and John T Hwang. Graph-accelerated non-intrusive polynomial chaos expansion using partially tensor-structured quadrature rules for uncertainty quantification. *Aerospace Science and Technology*, 155:109607, 2024.

[22] Mishal Thapa and Samy Missoum. Surrogate-based stochastic optimization of horizontal-axis wind turbine composite blades. *Structural and Multidisciplinary Optimization*, 65(2):41, 2022.

[23] Markus P Rumpfkeil. Optimizations under uncertainty using gradients, hessians, and surrogate models. *AIAA journal*, 51(2):444–451, 2013.

[24] Evin J Cramer, John E Dennis, Jr, Paul D Frank, Robert Michael Lewis, and Gregory R Shubin. Problem formulation for multidisciplinary optimization. *SIAM Journal on Optimization*, 4(4):754–776, 1994.

[25] Dongbin Xiu and George Em Karniadakis. The wiener–askey polynomial chaos for stochastic differential equations. *SIAM journal on scientific computing*, 24(2):619–644, 2002.

[26] Dongjin Lee and Sharif Rahman. Practical uncertainty quantification analysis involving statistically dependent random variables. *Applied Mathematical Modelling*, 84:324–356, 2020.

[27] John P Jasa, John T Hwang, and Joaquim RRA Martins. Open-source coupled aerostructural optimization using python. *Structural and Multidisciplinary Optimization*, 57(4):1815–1827, 2018.

[28] Mohamed Amine Bouhlel, Nathalie Bartoli, Abdelkader Otsmane, and Joseph Morlier. Improving kriging surrogates of high-dimensional design models by partial least squares dimension reduction. *Structural and Multidisciplinary Optimization*, 53(5):935–952, 2016.

[29] Mohamed Amine Bouhlel, John T Hwang, Nathalie Bartoli, Rémi Lafage, Joseph Morlier, and Joaquim RRA Martins. A python surrogate modeling framework with derivatives. *Advances in Engineering Software*, 135:102662, 2019.