



**HAL**  
open science

# **New Bayesian optimisation framework for robust multi-objective design: decoupling performance and uncertainty**

Hao Zhou, Zhongzheng Wang, Didier Lucor, Emilie Sauret

## ► To cite this version:

Hao Zhou, Zhongzheng Wang, Didier Lucor, Emilie Sauret. New Bayesian optimisation framework for robust multi-objective design: decoupling performance and uncertainty. *Reliability Engineering and System Safety*, 2026, pp.112359. <10.1016/j.ress.2026.112359>. <hal-05513238>

**HAL Id: hal-05513238**

**<https://hal.science/hal-05513238v1>**

Submitted on 16 Feb 2026

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-NC-ND 4.0 - Attribution - Non-commercial use - No Derivative Works - International License

# New Bayesian optimisation framework for robust multi-objective design: decoupling performance and uncertainty

Hao Zhou<sup>a</sup>, Zhongzheng Wang<sup>a</sup>, Didier Lucor<sup>b,\*</sup>, Emilie Sauret<sup>a,\*\*</sup>

<sup>a</sup>*School of Mechanical, Medical and Process Engineering, Faculty of Engineering, Queensland University of Technology, AU*

<sup>b</sup>*Université Paris-Saclay, CNRS, Laboratoire Interdisciplinaire des Sciences du Numérique, 91405 Orsay, France*

---

## Abstract

Robust optimisation is increasingly critical in engineering design due to the growing emphasis on robustness under uncertainty. However, existing methods often neglect explicit control of performance stability, particularly regarding the variance of objectives. To address this limitation, this work introduces a Variance-constrained Robust Bayesian Optimisation framework to efficiently perform multi-objective optimisation considering input uncertainty. A robust Gaussian process is employed to quantify input uncertainty by providing both the expected objective value and its associated uncertainty. These two aspects are then integrated using a variance-penalised scalarisation method, which leverages user-defined parameters to flexibly balance mean performance and variability, thereby accommodating different levels of stability tolerance. During optimisation, this scalarisation steers the acquisition function toward regions consistent with the desired stability preference; at reporting time, surrogate estimates are used to discard any designs that violate the uncertainty threshold, and the Pareto set is computed from the remaining stability-feasible points. The novelty lies in enabling a flexible, context-dependent design selection, while preserving the data-efficiency advantages of Bayesian Optimisation. The framework has been thoroughly tested on a series of synthetic benchmark problems, including a higher-dimensional (6D) case, demonstrating its effectiveness in handling varying uncertainty tolerances and substantiating its strong potential for real-world applications.

*Keywords:* Uncertainty Tolerance, Multi-objective Optimisation, Robust Optimisation, Stability-feasible Pareto fronts

---

---

\*corresponding: didier.lucor@cnrs.fr

\*\*corresponding: emilie.sauret@qut.edu.au

## Nomenclature

Notation	Description
<i>Multi-objective optimisation</i>	
$\mathbf{f}$	Vector of true black-box objective functions
$f_i$	$i$ -th objective function
$N$	Number of objectives
$x$	Scalar design variable
$M$	Dimension of design variables
$\mathbf{x}$	Vector of design variables (decision point)
$\mathcal{X}$	Feasible design space
$\mathcal{X}^*$	Pareto-optimal set
$\mathbb{F}$	Objective value space
$\mathcal{P}^*$	Pareto front in objective space
<i>Bayesian optimisation</i>	
$D^n$	Dataset of $n$ evaluated design points and their observations
$(i)$	Superscript denoting the $i$ -th data point in the dataset, e.g., $\mathbf{x}^{(i)}$ , $f^{(i)}$
$f^{\text{GP}}$	Gaussian process surrogate for objective $f$
$m(\cdot)$	Predictive mean function of the GP
$\sigma(\cdot)$	Predictive standard deviation (uncertainty) from the GP
$\hat{f}$	Predictive mean of the surrogate $f^{\text{GP}}$
$a(\cdot)$	Acquisition function
<i>Robust Bayesian optimisation</i>	
$\boldsymbol{\xi}$	Input perturbation/noise added to $\mathbf{x}$
$p_{\boldsymbol{\xi}}(\boldsymbol{\xi})$	Probability distribution over input perturbations
$\mathbf{x}$	Perturbed input: $\mathbf{x} + \boldsymbol{\xi}$
$J_i$	Robust objective: expectation over $f_i(\mathbf{x} + \boldsymbol{\xi})$
$\mathbf{J}$	Vector of robust objectives
$J_i^{\text{GP}}$	GP surrogate of robust objective $J_i$
$f^{\text{RGP}}$	Robust GP surrogate model (same as $J^{\text{GP}}$ )
$V_i$	Variance (robust uncertainty) of $f_i(\mathbf{x} + \boldsymbol{\xi})$
$\mathbf{V}$	Vector of variance
$\mathbf{p}\mathbf{v}$	Threshold for uncertainty in each objective
$\mathbf{y}_i$	Observation of $f_i(\mathbf{x} + \boldsymbol{\xi})$
$\mathcal{X}_{\text{cand}}$	Set of candidate points to be evaluated
$S$	Monte Carlo sample size for uncertainty estimation
$z/\theta$	Reference point in scalarisation (e.g., current best)
$\mathcal{H}$	Augmented robust objective
$\mathbf{H}$	Vector of augmented robust objectives

## 1. Introduction

In modern engineering practice, design problems are becoming increasingly complex and critical, particularly as performance demands and safety standards rise across industries. These problems often involve multiple conflicting objectives that must be optimised simultaneously. For instance, in battery thermal management, heat-exchanger design must balance peak temperature, heat removal, and pressure drop (hydraulic losses). It must do so within tight packaging and pumping-power limits, while keeping complexity and cost manageable.

Significant research has been conducted on Multi-Objective Optimisation (MOO) as well as optimisation under constraints [1, 2, 3]. To address such challenges, the field of MOO has developed extensively, offering methods to identify sets of trade-off solutions known as Pareto fronts. Conventional MOO algorithms, such as the Non-Dominated Sorting Genetic Algorithm II (NSGA-II) [4] and Strength Pareto Evolutionary Algorithm [5], have been widely applied in engineering design due to their ability to explore complex, high-dimensional design spaces and to approximate the true Pareto-optimal set.

However, conventional MOO methods typically assume deterministic conditions and do not explicitly account for input uncertainty, which limits their robustness when applied to real-world systems affected by noise and variability [6, 7]. For example, small geometrical variations in the surface feature of the interface materials in electronics applications (e.g., 1 micrometer increase in surface roughness) can lead to drastic change in the thermal resistance by a factor of 2 [8]. Therefore, addressing optimisation under input uncertainty has received increasing attention. Nevertheless, although robustness under uncertainty has been extensively explored in single-objective optimisation, its extension to the multi-objective setting has received comparatively less attention. In the broader optimisation literature, input uncertainty is often classified into robust formulations (distribution-free, set-based uncertainty) and stochastic formulations (distribution-based uncertainty) [9, 10]. In line with more recent work [11] that treats these categories under a unified stochastic viewpoint, we use the term robust in this paper in a distribution-based sense: the input uncertainty is described by a probability distribution, and robustness is measured through functionals of the resulting output distribution (in our case, the mean and the variance).

Two well-established paradigms are Robust-Design Optimisation (RDO) [12, 13] and Reliability-based Design Optimisation (RBDO) [14, 15, 16]. RDO expresses objectives (and possibly constraints) in terms of statistics such as expectation and variance, seeking designs with low sensitivity to input variation. RBDO focuses on the probability that constraints are respected; smaller admissible violation probabilities imply higher reliability. When surrogates are used—data-driven emulators of costly simulators that map inputs to outputs at a fraction of the evaluation cost—they

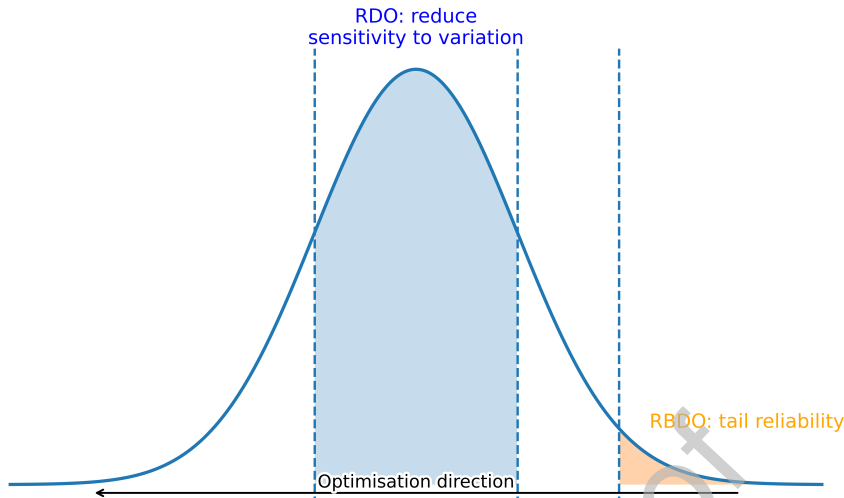


Figure 1: Illustration of Robust design optimisation (RDO) and Reliability-based Design Optimisation (RBDO).

41 replace repeated high-fidelity runs and, in the case of Gaussian-process models,  
 42 provide predictive uncertainty to guide exploration. Fig. 1 highlights the practical  
 43 distinction: we use the term RDO for formulations that optimise performance  
 44 measures based on response statistics such as the mean and, when appropriate, its  
 45 spread (e.g., variance) whereas RBDO explicitly limits the tail probability of vio-  
 46 lating a specified limit. In practice, multi-objective workflows often replace each  
 47 nominal objective by a robust statistic (e.g., mean or mean-variance) and, when  
 48 needed, monitor a simple reliability indicator (e.g., a probability of violation), with-  
 49 out committing to a single constrained formulation. Reliability-based Robust De-  
 50 sign Optimisation (RBRDO) [17, 18] combines these views by using expectations and  
 51 variances to describe performance while bounding failure probabilities for reliability.

52

53 A common way to adapt traditional MOO to uncertainty is to optimise statistics,  
 54 such as mean and maximum, in place of the nominal objectives. However, estimating  
 55 these statistics via repeated uncertainty sampling [19] (e.g., Exhausted NSGA-II  
 56 (ENSGA-II) [20]) is prohibitively costly, especially for expensive high-fidelity black-  
 57 box evaluations, motivating sample-efficient alternatives.

58

59 To address this challenge, Bayesian Optimisation (BO) [21] has emerged as a  
 60 powerful framework for sample-efficient optimisation of expensive black-box func-  
 61 tions. BO constructs a probabilistic surrogate model, typically a Gaussian Pro-  
 62 cess (GP) [22], that not only approximates the objective function but also quantifies  
 63 predictive uncertainty. Consequently, many studies have combined BO and Robust  
 Optimisation (RO) into Robust Bayesian Optimisation (RBO) [23, 24, 25]. These

64 methods incorporate input uncertainty directly into the BO loop by applying re-  
 65 peated sampling on GP surrogate rather than the original objective, which signifi-  
 66 cantly reduces computational cost. This leads to a Robust Gaussian Process (RGP)  
 67 that models the distribution of the perturbed response and provides uncertainty-  
 68 aware statistics (e.g., expectation or mean-variance). The RGP thus replaces the  
 69 nominal GP in guiding the acquisition function and downstream optimisation. Kudva  
 70 et al. [26] proposed flexibility-analysis algorithm that accommodated uncertain pa-  
 71 rameters and recourse variables across multiple constraints. Li et al. [27] utilised  
 72 a Bayes' risk in RGP, achieving robust uncertainty quantification to account for  
 73 the inherent stochasticity of additive manufacturing processes across different ge-  
 74 ometries and process parameters. Oliveira et al. [28] proposed handling both terms  
 75 via an upper confidence bound strategy in RBO, offering theoretical guarantees for  
 76 noisy-input settings and demonstrating competitive performance against other BO  
 77 approaches dealing with input uncertainty. Inatsu et al. [29] proposed a GP-based  
 78 MOBO that constructs high-probability bounding boxes for general risk measures  
 79 by mapping a high-probability credible region of the response, thereby screening out  
 80 extreme (unsafe) scenarios under input uncertainty.

81 Despite significant efforts have been made in the development of RBO, exist-  
 82 ing approaches still share a critical limitation: they focus on either the expectation  
 83 (Bayes' risk) or a worst-case scenario proxy (e.g., worst-case, Value-at-Risk (VaR)).  
 84 Expectation-based methods aim to optimise average performance under uncertainty,  
 85 which might inadvertently tolerate excessively large performance fluctuations around  
 86 the mean [30]. Conversely, worst-case approaches concentrate on ensuring robust  
 87 performance under extreme scenarios, potentially resulting in overly conservative  
 88 solutions that compromise average performance [31]. Crucially, neither approach ex-  
 89 plicitly controls the variance of the objective induced by input uncertainty [32, 33].  
 90 This overlooked variability can pose severe operational risks and ignore the impor-  
 91 tance of stability, especially in quality-sensitive applications, as significant fluctua-  
 92 tions in system response might frequently violate critical design constraints or safety  
 93 thresholds.

94 Inspired by the exploration-exploitation trade-off inherent in traditional acquisi-  
 95 tion functions such as Expected Improvement (EI), we propose a variance-penalised  
 96 scalarisation method that explicitly balances the expected performance and the un-  
 97 certainty in objectives caused by input perturbations. Through this approach, the  
 98 mean and standard deviation of each objective are combined into a single scalar  
 99 quantity, thereby enabling direct application of conventional acquisition strategies.  
 100 Recognising that the predictive variance of an RGP mixes two qualitatively differ-  
 101 ent sources—epistemic (model) uncertainty from limited data and aleatory (objec-  
 102 tive) uncertainty induced by input perturbations—we proposed a bi-loop Variance-  
 103 constrained Robust Bayesian Optimisation (VRBO) architecture with two parallel

branches updated within each iteration. The mixing occurs because a random input is propagated through a probabilistic model, so the resulting spread reflects both model ignorance and inherent input-driven variability and appears as a single “variance” signal to the optimiser. One branch refines the nominal surrogate to reduce epistemic error; the other targets aleatory variability using the robust surrogate and the variance-penalised scalarisation, steering the search toward stability-feasible regions. Treating these components separately is crucial: epistemic uncertainty should shrink with additional data, whereas aleatory uncertainty is irreducible and ought to shape "stability" preferences; if conflated, the optimiser may mistake undertrained areas for genuinely low-variance regions (or vice versa). This parallel design avoids that confounding and guides the optimisation toward robust solutions with predictable performance under input variability. Stability enters in two complementary forms: a soft influence within the acquisition, controlled by a tunable parameter  $\alpha$  that modulates the emphasis placed on uncertainty, and after optimisation, we filter the final Pareto front by a hard stability constraint, discarding any points that violate the user-specified tolerance  $\mathbf{p}_v$ .

By employing this framework, optimisation results can be controlled based on the specified tolerance to objective uncertainty. In this study, four benchmark cases are utilised to validate the proposed VRBO method. These cases are specifically selected according to categories defined by the relationship between the Pareto fronts obtained by non-robust and robust optimisation. Through the application of varying uncertainty tolerance levels within VRBO, the benchmark cases exhibit distinctly different Pareto front behaviours, demonstrating the method’s capability to generate flexible trade-offs between performance stability and optimality.

In summary, this study makes the following main contributions:

1. A variance-penalised scalarisation is introduced that explicitly integrates the input-uncertainty-induced mean and variance into the search, providing a user-tunable trade-off between expected performance and variability. This produces robust Pareto fronts tailored to specified uncertainty tolerances.
2. A bi-loop optimisation strategy VRBO is proposed to effectively balances global surrogate accuracy with local robustness requirements, ensuring stable optimisation performance under noisy inputs.
3. The effectiveness and versatility of the proposed VRBO method are demonstrated through an empirical study on five benchmarks. These results confirm that our approach consistently produces adaptive or competitive Pareto fronts compared to conventional optimisation methods.

The rest of this paper is organised as follows. In Section 2, we introduce the background knowledge about optimisation. Section 3 presents the structure of VRBO. The benchmark cases are discussed in Section 4, 5 and 6. Finally, we conclude and summarise our findings in Section 7.

## 144 2. Preliminaries

### 145 2.1. Multi-objective Optimisation

146 Many engineering problems require simultaneous optimisation of multiple con-  
 147 flicting objectives (with or without constraints). Given a vector-valued objective  
 148 function  $\mathbf{f} = (f_1, f_2, \dots, f_N)$ , the goal of MOO is to find solutions that balance  
 149 trade-offs between objectives [34].

150 Considering a MOO of minimizing each objective function, a candidate design  
 151 vector  $\mathbf{x}^* \in \mathcal{X} \subseteq \mathbb{R}^M$  is preferable to another  $\mathbf{x}$  one if:

$$152 \quad f_i(\mathbf{x}^*) \leq f_i(\mathbf{x}), \forall i \in \{1, \dots, N\}, \text{ and } \exists i \in \{1, \dots, N\} : f_i(\mathbf{x}^*) < f_i(\mathbf{x}). \quad (1)$$

153 In this case, these relations introduce the concept of dominance, denoted here as  
 154  $\mathbf{x} \prec \mathbf{x}^*$  (in the case of a minimisation problem) to indicate that  $\mathbf{f}(\mathbf{x}) \prec \mathbf{f}(\mathbf{x}^*)$ . In  
 155 MOO, a candidate  $\mathbf{x}^*$  is defined as Pareto optimal if:  $\nexists \mathbf{x} \in \mathcal{X}$  such that  $\mathbf{f}(\mathbf{x}^*) \prec$   
 156  $\mathbf{f}(\mathbf{x})$ .

157 The set of all Pareto optimal solutions is called the Pareto set  $\mathcal{X}^*$ :

$$\mathcal{X}^* := \{\mathbf{x}^* \in \mathcal{X} \mid \nexists \mathbf{f}(\mathbf{x}) \in \mathbb{F}_{\mathbf{f}} \text{ s.t. } \mathbf{f}(\mathbf{x}^*) \prec \mathbf{f}(\mathbf{x})\}, \quad \text{where } \mathbb{F}_{\mathbf{f}} := \{\mathbf{f}(\mathbf{x}) \mid \mathbf{x} \in \mathcal{X}\}, \quad (2)$$

158 and its image in the objective space is called the Pareto front  $\mathcal{P}^*$  [35], representing  
 159 optimal trade-offs where improving one objective inevitably worsens another. The  
 160 aim is thus to accurately approximate the Pareto front with as few expensive queries  
 161 to  $\mathbf{f}$  as possible. In this context, it makes sense to turn to efficient optimisation  
 162 techniques, for instance relying on objective functions model approximations.

### 163 2.2. Bayesian Optimisation framework

164 BO [21] is a sample-efficient framework for optimising expensive black-box func-  
 165 tions. It constructs a probabilistic surrogate – commonly an adaptive GP [36, 37] for  
 166 which we briefly introduce the key ingredients here – to model the objective func-  
 167 tion with a continuous representation of  $f(\mathbf{x})$  from a discrete dataset, and uses an  
 168 acquisition function to guide the search toward promising regions of the input space  
 169 in order to refine and improve the surrogate.

170 Given a dataset  $D^n = \{(\mathbf{x}^{(i)}, f^{(i)})\}_{i=1}^n$ , a GP defines a prior over functions spec-  
 171 ified by a mean function  $m(\mathbf{x})$  and a covariance kernel  $k(\mathbf{x}, \mathbf{x}')$ , representing the  
 172 independence of function values at different locations. Conditioned on the data, the  
 173 GP posterior at any test input  $\mathbf{x}_*$  is:

$$f^{\text{GP}}(\mathbf{x}_*) | D^n \sim \mathcal{N}(m_{f^{\text{GP}}}(\mathbf{x}_*), \sigma_{f^{\text{GP}}}^2(\mathbf{x}_*)), \quad (3)$$

174 where the posterior mean and the variance depend on the choice of the kernel function  
 175 and the observed data. On the diagonal,  $k(\mathbf{x}, \mathbf{x})$  gives the prior marginal variance

176 at  $\mathbf{x}$ ; after conditioning, the posterior predictive variance  $\sigma_{f_{\text{GP}}}^2(\mathbf{x}_* | D^n)$  quantifies  
 177 local uncertainty (and typically shrinks near observed points). And the surrogate  
 178 prediction is the GP posterior mean,

$$\hat{f}(\mathbf{x}_*) = m_{f_{\text{GP}}}(\mathbf{x}_* | D^n). \quad (4)$$

179 The acquisition function  $a(\mathbf{x})$  quantifies the utility of sampling a new point  $\mathbf{x}$ , balanc-  
 180 ing exploration (high variance) and exploitation (low predicted mean). The choice  
 181 of this acquisition function is crucial, as one easily understands that it is better to  
 182 refine the surrogate only if necessary, and possibly in the close neighbourhood of  
 183 the optimal design parameters region. Common choices include EI, Probability of  
 184 Improvement, and Upper Confidence Bound [38, 39]. At each iteration  $k$ , the next  
 185 sample is chosen as:

$$\mathbf{x}^{(k+1)} = \arg \max_{\mathbf{x}} a(f^{\text{GP}}(\mathbf{x})). \quad (5)$$

186 In Multi-Objective Bayesian Optimisation (MOBO) [40, 41], a simple choice is  
 187 to rely on separate GP surrogates, one for each objective function  $f_i$ , and acquisition  
 188 functions such as Expected Hypervolume Improvement (EHVI), are used to suggest  
 189 candidates that efficiently improve the Pareto front.

### 190 2.3. Accounting for uncertainties

191 In many real-world systems, the exact values of actual input parameters are not  
 192 precisely known. This uncertainty may arise from factors such as manufacturing and  
 193 assembly tolerances, or inaccuracies in material properties. It is therefore essential  
 194 to incorporate these parameter uncertainties into the optimisation process to ensure  
 195 the resulting design is robust.

196 Such uncertainties introduce non-determinism into the system's response and  
 197 constraints, making the optimisation process considerably more challenging. To ad-  
 198 dress this, one must first define a deterministic quantity of interest to optimize. This  
 199 typically involves statistical measures – such as moments (e.g., mean or variance),  
 200 probabilities, or specific events like worst-case scenarios – related to the system's  
 201 performance.

202 Another important consideration is identifying which parameters are affected by  
 203 uncertainty. In some cases, uncertainties pertain to environmental or operational  
 204 conditions that are independent of the design variables  $\mathbf{x}$ . However, in the scenario  
 205 considered in this study, the uncertainties directly impact the design variables them-  
 206 selves. As a result, the actual system input takes the form  $\mathbf{x} + \boldsymbol{\xi}$ , where  $\boldsymbol{\xi}$  is a known  
 207 zero-mean random vector representing manufacturing variabilities. This deviation  
 208 from the nominal design must be accounted for in order to achieve a robust and  
 209 reliable solution.

210 To account for such perturbations, we consider not only the nominal objec-  
 211 tive value  $\mathbf{f}(\mathbf{x})$ , but its statistical response under uncertainty—specifically, the mean

212 (Bayes' risk) and variance—leading to the formulation of robust and risk-aware opti-  
213 misation strategies [42].

### 214 3. Variance-constrained Robust Bayesian Optimisation

#### 215 3.1. Optimisation under uncertainties formulation

216 We consider a deterministic design  $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^M$  and input perturbations  $\boldsymbol{\xi} \sim p_{\boldsymbol{\xi}}$ .  
217 Each objective is evaluated as  $f_i(\mathbf{x} + \boldsymbol{\xi})$ . To summarise uncertainty effects we use:

$$\begin{aligned} J_i(\mathbf{x}) &= \mathbb{E}_{\boldsymbol{\xi} \sim p_{\boldsymbol{\xi}}}[f_i(\mathbf{x} + \boldsymbol{\xi})], & (\text{robust / Bayes' risk}) \\ V_i(\mathbf{x}) &= \mathbb{V}_{\boldsymbol{\xi} \sim p_{\boldsymbol{\xi}}}[f_i(\mathbf{x} + \boldsymbol{\xi})], & (\text{stability / uncertainty}) \end{aligned} \quad (6)$$

218 In our setting, “stability” is operationalised as compliance with a user-specified  
219 tolerance on input-induced variability, i.e.  $V_i(\mathbf{x}) \leq p_{V_i}$ . We therefore search in the  
220  $\mathbf{J}$ -space while enforcing the hard variance bound at decision time.

221 *Robust multi-objective design optimisation.* Under the MOO background, we aim  
222 to find Pareto-optimal solutions that minimise the Bayes' risk vector  $\mathbf{J}(\mathbf{x})$  while  
223 satisfying user-defined tolerance  $\mathbf{p}_V$  on the variance vector  $\mathbf{V}(\mathbf{x})$ :

$$\begin{aligned} \min_{\mathbf{x} \in \mathcal{X}} \quad & \mathbf{J}(\mathbf{x}) = (J_1(\mathbf{x}), \dots, J_N(\mathbf{x})), \\ \text{s.t.} \quad & \mathbf{V}(\mathbf{x}) = (V_1(\mathbf{x}), \dots, V_N(\mathbf{x})) \leq \mathbf{p}_V, \end{aligned} \quad (7)$$

224 where  $\mathcal{X} \subseteq \mathbb{R}^M$  is the feasible domain, and  $\mathbf{p}_V \in \mathbb{R}_+^N$  is a user-defined vector  
225 of variance thresholds. Note that any maximisation objective can be recast as a  
226 minimisation by negating it.

227 To solve this stochastic optimisation efficiently under limited data budgets, we  
228 require surrogate models that can account for the propagation of input uncertainty.

#### 229 3.2. Robust Gaussian Process surrogate

230 In order to model each objective  $f_i$  under input uncertainty given a dataset,  
231 we construct a RGP surrogate. This model approximates the Bayes' risk  $J_i(\mathbf{x}) =$   
232  $\mathbb{E}_{p_{\boldsymbol{\xi}}}[f_i(\mathbf{x} + \boldsymbol{\xi})]$  while bounding its uncertainty using a modified GP framework.

233 *Kernel expectation formulation.* Let  $k_{f_i}(\mathbf{x}, \mathbf{x}')$  denotes the kernel of the nominal GP  
234 surrogate  $f_i^{\text{GP}}$ , and  $K_i = K_{f_i} + \sigma_{\varepsilon}^2 I$  the posterior covariance matrix based on the  
235 training outputs  $\mathbf{y}_i$ . The robust surrogate  $J_i^{\text{GP}} \sim \mathcal{GP}(m_{J_i}, \text{Cov}_{J_i})$  is then given  
236 by [22]:

$$m_{J_i}(\mathbf{x}) = k_{J_i, f_i}(\mathbf{x})^\top K_i^{-1} \mathbf{y}_i, \quad (8)$$

$$\text{Cov}_{J_i}(\mathbf{x}, \mathbf{x}') = k_{J_i}(\mathbf{x}, \mathbf{x}') - k_{J_i, f_i}(\mathbf{x})^\top K_i^{-1} k_{J_i, f_i}(\mathbf{x}'), \quad (9)$$

237 where the kernel expectations are defined as:

$$k_{J_i, f_i}(\mathbf{x}) = \int k_{f_i}(\mathbf{x} + \boldsymbol{\xi}) p(\boldsymbol{\xi}) d\boldsymbol{\xi}, \quad (10)$$

$$k_{J_i}(\mathbf{x}, \mathbf{x}') = \iint k_{f_i}(\mathbf{x} + \boldsymbol{\xi}, \mathbf{x}' + \boldsymbol{\xi}') p(\boldsymbol{\xi}) p(\boldsymbol{\xi}') d\boldsymbol{\xi} d\boldsymbol{\xi}'. \quad (11)$$

238 For squared–exponential kernels and Gaussian perturbations, the integrals defin-  
 239 ing the robust objective can be computed analytically. Otherwise, we approximate  
 240 them using Monte Carlo integration [24]. Specifically, for each candidate input  $\mathbf{x}$ ,  
 241 a set of perturbations  $\{\boldsymbol{\xi}^{(s)}\}_{s=1}^S$  is drawn from the input noise distribution  $p_{\boldsymbol{\xi}}$ , and  
 242 the corresponding perturbed inputs  $\mathbf{x} + \boldsymbol{\xi}^{(s)}$  are evaluated using the nominal GP  
 243 surrogate  $f^{\text{GP}}$ . The empirical mean and variance of these surrogate outputs are then  
 244 used to approximate the robust mean  $m_J(\mathbf{x})$  and variance  $\text{Cov}_J(\mathbf{x})$ . These statistics  
 245 are subsequently used to construct the robust surrogate  $J^{\text{GP}}$  with heteroscedastic  
 246 noise (the noise level is various across different input).

247 *Practical construction.* The following procedure builds a RGP surrogate by re–using  
 248 the nominal GP and propagating noise through the model:

---

**Algorithm 1:** Constructing a Robust Gaussian Process under input uncer-  
 tainty (here one–dimensional for mono–objective case)

---

**Input** : Dataset  $D^n = \{(\mathbf{x}^{(i)}, y^{(i)})\}_{i=1}^n$ ;  
 Perturbation distribution function  $p_{\boldsymbol{\xi}}(\boldsymbol{\xi})$ ;  
 Candidate set  $\mathcal{X}_{\text{cand}}$ ;  
 Sample count for Monte Carlo  $S$

**Output:** Robust surrogate  $J^{\text{GP}}$

- 249 1 Train nominal GP surrogate  $f^{\text{GP}}$  on  $D^n$ ;
  - 2 **foreach**  $\mathbf{x} \in \mathcal{X}_{\text{cand}}$  **do**
  - 3     Draw  $S$  samples  $\{\boldsymbol{\xi}^{(s)}\}_{s=1}^S \sim p_{\boldsymbol{\xi}}(\boldsymbol{\xi})$ ;
  - 4     Predict with surrogate  $y^{(s)} \leftarrow f^{\text{GP}}(\mathbf{x} + \boldsymbol{\xi}^{(s)})$ ;
  - 5     Estimate;
  - 6      $m_J(\mathbf{x}) \leftarrow \frac{1}{S} \sum_{s=1}^S y^{(s)}$ ;
  - 7      $\text{Cov}_J(\mathbf{x}) \leftarrow \mathbb{V}[y^{(1)}, \dots, y^{(S)}]$ ;
  - 8 Fit  $J^{\text{GP}}$  on  $\{(\mathbf{x}, m_J(\mathbf{x}))\}$  with heteroscedastic noise  $\text{Cov}_J(\mathbf{x})$ ;
- 

250 In all experiments, the robust statistics are approximated by Monte Carlo with  
 251 a fixed sample size of  $S = 200$  per design point [43], while the candidate set  $\mathcal{X}_{\text{cand}}$   
 252 is generated by Sobol sequences with a dimension–dependent size. The resulting  
 253 surrogate captures both the mean Bayes’ risk  $m_{J_n}(\mathbf{x})$  and its epistemic uncertainty  
 254  $\sigma_{J_n}^2(\mathbf{x}) = \text{Cov}_{J_n}(\mathbf{x}, \mathbf{x})$ , enabling the construction of variance–penalised acquisition  
 255 functions discussed in Section 3.4.

### 256 3.3. Existing acquisition strategies for (robust) MOBO

257 Before introducing our variance-penalised scalarisation, we briefly situate it within  
 258 existing acquisition strategies for multi-objective and RBO. In standard MOBO, a  
 259 common approach is to build one GP per objective and either (i) scalarise the vector  
 260 of objectives into a single surrogate (e.g. ParEGO, which combines random weights  
 261 with EI on a Tchebycheff-type scalarisation), or (ii) use Pareto-based criteria such as  
 262 EHVI and related hypervolume-based acquisitions, which directly target expansion  
 263 of the dominated region in the objective space.

264 When input uncertainty is present, many “robust MOBO” methods follow a  
 265 similar pattern: they first define robust statistics of each objective—typically ex-  
 266 pectations, mean-variance combinations, or risk measures—and then apply standard  
 267 single- or multi-objective acquisitions (e.g. EI/Upper Confidence Boundary (UCB)  
 268 or EHVI) to these transformed objectives. In other words, robustness enters through  
 269 the choice of response functional, while the acquisition itself remains conventional.

270 Our framework follows this general pattern but targets a different robustness  
 271 notion. In this work, We also adopt EHVI as the underlying MOBO acquisition,  
 272 yet we construct a dedicated scalar score  $H(\mu, \sigma)$  that combines the robust mean  
 273 and its input-induced variability in a tunable way, and we complement this soft,  
 274 acquisition-level trade-off with a hard variance filter at selection time. The next  
 275 subsection details this variance-penalised scalarisation.

### 276 3.4. Variance-penalised Scalarisation Algorithm

277 After constructing  $J_i^{\text{GP}}$ , we build a soft, risk-aware score that balances the robust  
 278 mean  $m_{J_n}(\mathbf{x})$  and the input-induced variability measured by  $\sigma_{J_n}^2(\mathbf{x})$ . The Conditional  
 279 Value-at-Risk (CVaR) term acts as a tail-severity penalty around a reference level,  
 280 while the  $\ln(1 + \sigma_i)$  term softly discourages large variance. This scalarisation guides  
 281 the search; final stability is enforced by the hard variance bound  $V_i(\mathbf{x}) \leq p_{V,i}$  when  
 282 reporting the Pareto set.

283 *Threshold-CVaR.* Let the best observed value in the dataset  $D^t$  for a minimisation  
 284 problem be  $\theta_i = \min_{y_i \in D^t} y_i$ . For any candidate  $\mathbf{x}$  with predictive mean  $\mu_i(\mathbf{x}) =$   
 285  $m_{J_i}(\mathbf{x})$  and standard deviation  $\sigma_i(\mathbf{x}) = \sqrt{\mathbb{V}[f_i(\mathbf{x} + \boldsymbol{\xi})]}$ , we consider the conditional  
 286 expectation of  $f_i(\mathbf{x} + \boldsymbol{\xi})$  above the target level  $\theta_i$ :

$$\text{CVaR}_{\theta_i}(\mathbf{x}) = \mathbb{E}[f_i(\mathbf{x} + \boldsymbol{\xi}) \mid f_i(\mathbf{x} + \boldsymbol{\xi}) \geq \theta_i].$$

287 Under the usual Gaussian-process assumption  $f_i(\mathbf{x} + \boldsymbol{\xi}) \sim \mathcal{N}(\mu_i(\mathbf{x}), \sigma_i^2(\mathbf{x}))$ , this  
 288 quantity has the closed form

$$\text{CVaR}_{\theta_i}(\mathbf{x}) = \mu_i(\mathbf{x}) + \sigma_i(\mathbf{x}) \frac{\phi(z_i(\mathbf{x}))}{1 - \Phi(z_i(\mathbf{x}))}, \quad z_i(\mathbf{x}) = \frac{\theta_i - \mu_i(\mathbf{x})}{\sigma_i(\mathbf{x})}, \quad (12)$$

289 where  $\phi$  and  $\Phi$  denote the standard normal density and cumulative distribution  
 290 function, respectively. In other words,  $z_i(\mathbf{x})$  is the standardised distance between  
 291 the target  $\theta_i$  and the predictive mean.

292 This quantity is closely related to the classical CVaR [32, 44]. In the standard  
 293 definition,  $\text{CVaR}_p$  at level  $p \in (0, 1)$  is the expected value of the loss conditional on  
 294 the loss exceeding its  $p$ -quantile. If, instead of fixing  $\theta_i$  to the current best observed  
 295 value, we chose  $\theta_i$  so that  $\mathbb{P}(f_i(\mathbf{x} + \boldsymbol{\xi}) \leq \theta_i) = p$ , then  $\text{CVaR}_{\theta_i}(\mathbf{x})$  would coincide with  
 296 the usual  $\text{CVaR}_p$ . In this work, however, we fix  $\theta_i$  as the current best observed value  
 297 and use  $\text{CVaR}_{\theta_i}(\mathbf{x})$  as a threshold-based tail risk measure: candidates for which the  
 298 upper tail above  $\theta_i$  has smaller conditional mean are preferred, as they are less likely  
 299 to produce severely poor outcomes relative to the current best performance.

300 *Variance-penalised scalarisation.* We then propose a modification and construct:

$$\mathcal{H}_i(\mathbf{x}) = \text{CVaR}_{\theta_i}(\mathbf{x}) + \alpha \ln[1 + \sigma_i(\mathbf{x})], \quad (13)$$

301 where  $\alpha \geq 0$  is a trade-off parameter that modulates the trade-off between tail-  
 302 risk aversion and exploration. Data samples for the next evaluation are chosen via  
 303  $\mathbf{x}^{(t+1)} = \arg \min_{\mathbf{x} \in \mathcal{X}_{\text{cand}}} \mathcal{H}_i(\mathbf{x})$ .

304 *Interpretation.* In order to summarise, we can say that the first term in Eq. (13)  
 305 captures a target-based exploitation strategy. Rather than simply being risk-averse,  
 306 it evaluates the distribution relative to a specific threshold; minimising this term  
 307 penalises any region whose performance is expected to be worse than the current  
 308 best. This ensures the acquisition function actively avoids areas where the objective  
 309 is likely to underperform compared to the best solution found so far. The second  
 310 log-variance term serves as a complementary exploration incentive that grows with  
 311 uncertainty. It acts as an exploration bonus: points with higher uncertainty get a  
 312 higher value, but not excessively so, due to the logarithm. Finally, the +1 term is  
 313 included to avoid singularities when the standard deviation tends to zero ( $\sigma \rightarrow 0 \Rightarrow$   
 314  $\ln(1 + \sigma) \approx \sigma$ ). This yields a search strategy that actively avoids underperformance  
 315 relative to the current best while remaining explorative, making it suitable for robust  
 316 optimisation under input uncertainty as we will see in the applications section.

### 317 3.5. Effect of $\alpha$

318 To visualise how the variance-penalised objective  $\mathcal{H}$  trades off mean performance  
 319 and predictive uncertainty, we evaluate  $\mathcal{H}$  on a grid  $\mu \in [0, 1]$ ,  $\sigma \in [0, 0.5]$  with the  
 320 CVaR reference fixed at  $\theta = 0.2$ . Fig. 2 shows maps for four settings of the trade-off  
 321 parameter  $\alpha$ . When  $\alpha = 0$  (Fig. 2 (a)),  $\mathcal{H}$  reduces to the CVaR term: isolines are  
 322 mostly vertical over much of the domain (mean dominates), but they can flatten for  
 323 very small  $\mu$  and become diagonal where  $\sigma$  is large. With a small penalty ( $\alpha = 0.1$ ,  
 324 Fig. 2 (b)) the contours begin to bend, reflecting a mild preference for smaller  $\sigma$ . At

325  $\alpha = 1$  (Fig. 2 (c)) the contours tilt towards  $\approx 45^\circ$ , indicating a balanced emphasis  
 326 on reducing both tail loss and variance. For a strongly risk-averse choice ( $\alpha = 10$ ,  
 327 Fig. 2 (d)), the variance penalty dominates: contours become almost horizontal  
 328 and designs with larger  $\sigma$  are penalised irrespective of  $\mu$ . Overall, as  $\alpha$  increases,  
 329 the acquisition places more weight on low variance, so the search shifts from mean-  
 330 seeking (small  $\alpha$ ) to variance-averse sampling that concentrates evaluations in low- $\sigma$   
 331 regions.




Figure 2: Contour maps of the scalar objective  $\mathcal{H}(\mu, \sigma)$  for (a)  $\alpha = 0.$ , (b)  $\alpha = 0.1$ , (c)  $\alpha = 1.0$ ,  
 (d)  $\alpha = 10$ . The horizontal axis is the predictive mean  $\mu$ , the vertical axis the standard  
 deviation  $\sigma$ . Warmer colours indicate a higher (worse) objective value.

### 3.6. Variance-constrained Robust Bayesian Optimisation Framework

332 In a RGP, the predictive variance  $\sigma_{J_i}^2(\mathbf{x})$  will be affected by two distinct com-  
 333 ponents: the model uncertainty that reflects the surrogate’s lack of data and the  
 334 objective uncertainty that originates from the random input perturbations  $\boldsymbol{\xi}$ . An  
 335 illustrative example is shown in Fig. 3. Panel (a) displays the nominal GP for  $f(x)$ :  
 336 the blue curve is the posterior mean and the light-blue ribbon is the 95% epistemic  
 337 band. Panel (b) shows the robust surrogate  $J(x) = \mathbb{E}[f(x + \boldsymbol{\xi})]$  obtained by prop-  
 338 agating input noise through the GP: the orange curve is the robust mean and the  
 339 hatched ribbon is the 95% total band (epistemic  $\oplus$  aleatory). This highlights that  $J$   
 340 carries both model and input-induced uncertainty.

341 The VRBO framework therefore employs two nested optimisation loops, each  
 342 targeting one of these sources, as shown in Fig. 4.

344 *MO loop – Surrogate-model refinement.* The MO loop maintains nominal GP sur-  
 345 rogates  $\text{GP}_{f_i}$  for every objective and selects new points that minimise model un-  
 346 certainty. Typical criteria include maximising the GP posterior variance or using  
 347 information-based acquisition functions; the goal is to densify data where the surro-  
 348 gate is least confident, thereby shrinking epistemic error.




Figure 3: Uncertainty propagation from the nominal GP to the robust surrogate. (a) Nominal  $f(x)$ : GP mean (blue) with 95% epistemic band. (b) Robust  $J(x) = \mathbb{E}[f(x + \xi)]$ : mean (orange) with 95% total band (epistemic  $\oplus$  aleatory) obtained by propagating input perturbations through the GP.

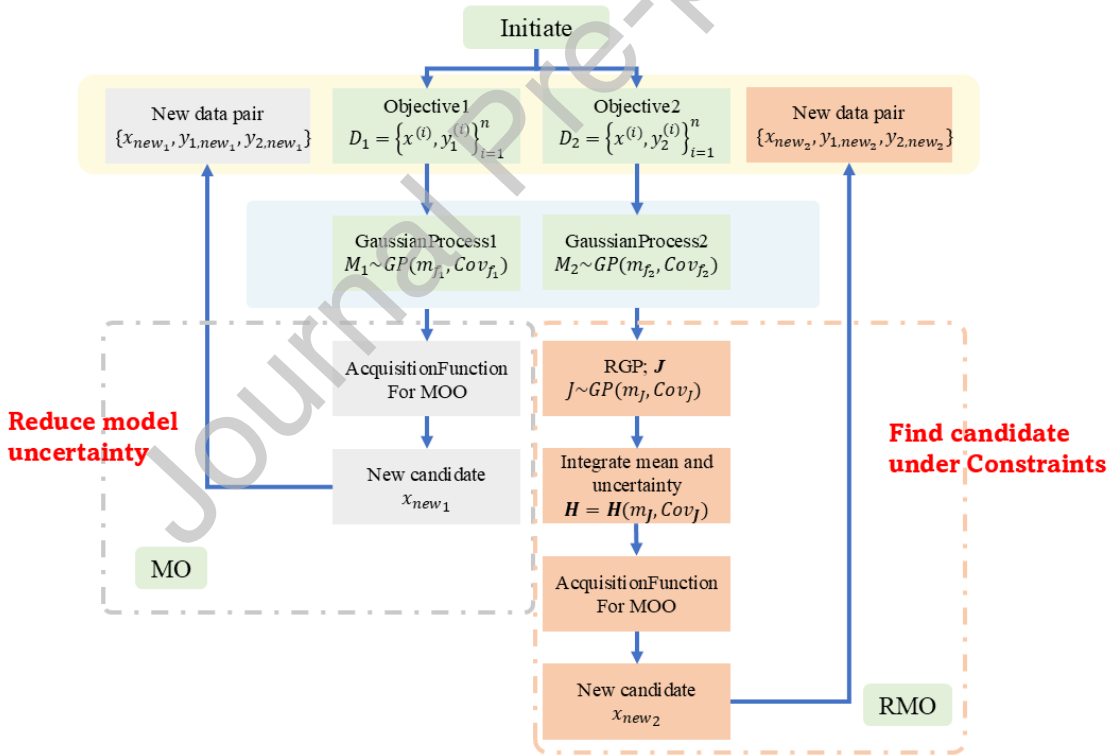


Figure 4: Bi-loop structure of the proposed VRBO framework. The MO loop (left) updates nominal GP models to reduce model uncertainty, while the RMO loop (right) uses a RGP surrogate and a variance-penalised scalarisation function  $H$  to find candidates under input uncertainty.

349 *RMO loop – Robust candidate selection.* The RMO loop builds RGP surrogates  $J_i^{\text{GP}}$   
 350 (Section 3.2) and employs the variance-penalised robust acquisition  $\mathcal{H}_i(\mathbf{x})$  in Eq. (13).  
 351 Optimisation of  $\mathcal{H}_i$  focuses on controlling objective uncertainty introduced by input  
 352 noise, producing designs whose Bayes–risk means and variances approximately satisfy  
 353 the user–set tolerance  $\mathbf{p}_v$ .

354 For the multi–objective case, the robust scores are aggregated—e.g. with EHVI,  
 355 ParEGO, or an  $\varepsilon$ –constraint strategy—and the next candidate is

$$\mathbf{x}^{(t+1)} = \arg \min_{\mathbf{x} \in \mathcal{X}_{\text{cand}}} a(\mathbf{H}(\mathbf{x})),$$

356 where  $a(\cdot)$  denotes the chosen acquisition function for MOO,  $\mathbf{H} = (\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_N)$ .  
 357 In this work, EHVI is adopted as the acquisition function for MOO, since EHVI di-  
 358 rectly maximises expected hypervolume improvement and tends to provide a more  
 359 balanced coverage of the Pareto front.

360 *Iterative update.* After evaluating the black–box model at  $\mathbf{x}^{(t+1)}$  and recording ob-  
 361 servation  $\mathbf{y}^{(t+1)}$ , the data set  $D^t$  is augmented and both loops repeat until a stopping  
 362 criterion (maximum calls or Pareto convergence) is reached.

363 By decoupling model and objective uncertainties, this nested framework simul-  
 364 taneously (i) refines global surrogate fidelity (MO) and (ii) steers sampling toward  
 365 designs that remain performant under input noise (RMO), achieving sample–efficient  
 366 and uncertainty–aware optimisation. Overall, the algorithm for the entire optimisa-  
 367 tion process is presented below in Algorithm 2.

---

**Algorithm 2:** Variance-constrained Robust Bayesian Optimisation for MOO
 

---

**Input** : Initial dataset  $D^t = \{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^t$ ;  
 Perturbation distribution  $p(\boldsymbol{\xi})$ ;  
 Feasible domain  $\mathcal{U}$ , variance-tolerance  $\mathbf{p}_V$ ;  
 Trade-off coefficient  $\alpha$ , evaluation budget  $T$ ;  
 MOO acquisition  $a_{\text{MOO}}(\cdot)$ , candidate pool  $\mathcal{X}_{\text{cand}}$   
**Output:** Final surrogates  $\hat{\mathbf{f}}(\mathbf{x})$

---

```

1 Initialise nominal  $\text{GP}_{\mathbf{f}} = (\text{GP}_{f_1}, \text{GP}_{f_2}, \dots, \text{GP}_{f_n})$  with  $D^t$ ;
2 while  $t < T$  do // single loop, dual branches
3   Retrain  $\text{GP}_{\mathbf{f}}$  on  $D^t$ ;
   ; // MO loop
4    $\mathbf{x}_{\text{MO}} = \arg \max_{\mathbf{x} \in \mathcal{X}_{\text{cand}}} [\text{GP}_{\mathbf{f}}(\mathbf{x})]$  ;
   ; // RMO loop
5   Construct robust surrogates  $\text{GP}_{\mathbf{J}} \sim \mathcal{GP}(m_{\mathbf{J}}, \text{Cov}_{\mathbf{J}})$  via Alg. 1 using  $\text{GP}_{\mathbf{f}}$ 
   and  $p(\boldsymbol{\xi})$ ;
6   foreach  $n = 1, \dots, N$  do
7      $\theta_n \leftarrow \min_{\mathbf{x} \in D^t} y_n(\mathbf{x})$ ;
8     foreach  $\mathbf{x} \in \mathcal{X}_{\text{cand}}$  do
9        $\mathbf{H}(\mathbf{x}) = \text{CVaR}_{\theta}(\mathbf{x}) + \alpha \ln(1 + \sigma_{\mathbf{J}}(\mathbf{x}))$  ;
10     $\mathbf{x}_{\text{RMO}} = \arg \max_{\mathbf{x} \in \mathcal{X}_{\text{cand}}} a_{\text{MOO}}(\mathbf{H}(\mathbf{x}))$ ;
11    Evaluate black box at  $\{\mathbf{x}_{\text{MO}}, \mathbf{x}_{\text{RMO}}\}$  and append results to  $D^t$  to obtain
     $D^{t+2}$ ;
12     $t \leftarrow t + 2$ 
13 Refit surrogate model  $\hat{\mathbf{f}}(\mathbf{x})$  based on final dataset  $D^t \mathcal{P}_{\mathbf{H}}$ ;
14 return  $\hat{\mathbf{f}}(\mathbf{x})$ ;

```

---

### 3.7. Stability-filtered robust Pareto front

Given the final surrogates  $\hat{\mathbf{f}}(\cdot)$  by VRBO, we estimate for each candidate  $\mathbf{x} \in \mathcal{X}$  the robust mean  $\boldsymbol{\mu}(\mathbf{x}) \approx \mathbf{J}(\mathbf{x})$  and uncertainty  $\boldsymbol{\sigma}(\mathbf{x}) \approx \sqrt{\mathbf{V}(\mathbf{x})}$  by propagating input noise through the surrogate (RGP in closed form, or MC otherwise). We then apply a hard stability filter  $\boldsymbol{\sigma}(\mathbf{x}) \leq \mathbf{p}_V$  and perform non-dominated sorting, using the implementation provided in BoTorch [45], on  $\boldsymbol{\mu}(\mathbf{x})$ . The resulting set is our reported robust Pareto front (RMO). The process is presented in Algorithm 3. In our pipeline, the parameter  $\alpha$  primarily affects the sampling policy during optimisation. After the data have been collected, however, we refit the surrogate on all observations and perform a separate global scan over the design space, applying the variance filter

379  $\sigma(\mathbf{x}) \leq \mathbf{p}_V$  and non-dominated sorting on  $\boldsymbol{\mu}(\mathbf{x})$ . Thus,  $\alpha$  mainly controls where we  
 380 sample during the optimisation loop, while the final robust Pareto set is determined  
 381 in a post-hoc manner by the converged surrogate and the chosen tolerance  $\mathbf{p}_V$ .

---

**Algorithm 3:** Robust evaluation & variance-filtered Pareto selection

---

**Input:**  $\hat{\mathbf{f}}(\cdot)$ : trained surrogate(s);

$\mathcal{X} \subset \mathbb{R}^d$ : candidate designs;

$p_\xi$ : input-perturbation law;

$S$ : Monte-Carlo samples per design;

$\mathbf{p}_V$ : per-objective uncertainty thresholds

**Output:**  $\mathcal{P}_{\text{RMO}}$ : stability-feasible robust Pareto set

1 **foreach**  $\mathbf{x} \in \mathcal{X}$  **do**

2     Draw  $\{\boldsymbol{\xi}^{(s)}\}_{s=1}^S \sim p_\xi$ ;

3      $\mathbf{y}^{(s)} \leftarrow \hat{\mathbf{f}}(\mathbf{x} + \boldsymbol{\xi}^{(s)})$ ,  $s = 1, \dots, S$ ;

4      $\boldsymbol{\mu}(\mathbf{x}) \leftarrow \frac{1}{S} \sum_{s=1}^S \mathbf{y}^{(s)}$ ;

5      $\boldsymbol{\sigma}(\mathbf{x}) \leftarrow \sqrt{\frac{1}{S-1} \sum_{s=1}^S (\mathbf{y}^{(s)} - \boldsymbol{\mu}(\mathbf{x}))^2}$ ;

6  $\mathcal{X}_{\text{feas}} \leftarrow \{\mathbf{x} \in \mathcal{X} \mid \boldsymbol{\sigma}(\mathbf{x}) \preceq \mathbf{p}_V\}$ ;

7  $\mathcal{P}_{\text{RMO}} \leftarrow \text{NONDOMINATEDSORT}(\mathcal{X}_{\text{feas}}, \text{objectives} = \boldsymbol{\mu})$ ;

8 **return**  $\mathcal{P}_{\text{RMO}}$ ;

---

## 383 4. Noise-free Observations

384 In practical engineering optimisation, data typically come from two sources: high-  
 385 fidelity simulations and physical experiments. In the simulation regime, the com-  
 386 manded design  $\mathbf{x}$  is realised exactly and the outputs are deterministic, so the training  
 387 data are noise-free. Input uncertainty is therefore not part of the data generation;  
 388 instead, it is accounted for at evaluation time by propagating perturbations through  
 389 the RGP (Algorithm 2). This setup allows us to isolate the intrinsic performance of  
 390 the proposed framework under ideal observations, which is validated by three bench-  
 391 mark cases in this section. The complementary experimental regime—inaccurate  
 392 in observations—is addressed later in Section 5, where we demonstrate our method  
 393 under noisy observations.

### 394 4.1. Detailed Analysis: MDTP3

395 This subsection presents a comprehensive investigation of the MDTP3 problem,  
 396 which serves as the primary benchmark for assessing the proposed framework.

397 *4.1.1. Ground Truth Analysis*

398 The MDTP3 problem is defined over a two-dimensional domain  $(x_1, x_2) \in [0, 1]^2$   
 399 and employs trigonometric or polynomial components to generate complex non-  
 400 dominated fronts. By adjusting problem parameters or the input domain, the degree  
 401 of conflict between objectives can be controlled. In this work, we adopt the following  
 402 definition:

$$\begin{aligned} f_1(x_1, x_2) &= x_1, \\ f_2(x_1, x_2) &= 1 - 0.9 \exp\left[-\left(\frac{x_2 - 0.8}{0.1}\right)^2\right] - 1.3 \exp\left[-\left(\frac{x_2 - 0.3}{0.03}\right)^2\right] - x_1. \end{aligned} \quad (14)$$

403 The optimisation task aims to minimise both objectives:

$$\min_{(x_1, x_2) \in [0, 1]^2} [f_1(x_1, x_2), f_2(x_1, x_2)].$$

404 To evaluate the robustness of the proposed framework, Gaussian input noise with  
 405 standard deviation  $\sigma_{x_1} = \sigma_{x_2} = 0.02$  is introduced to both design variables. This  
 406 noise level is chosen because it produces a ground truth mean and standard deviation  
 407 landscape with conflicting characteristics, as illustrated in Fig. 5, generated by high-  
 408 dimension grid search combined with Monte Carlo evaluation. At  $x_2 = 0.3$  the mean  
 409 performance is better than at  $x_2 = 0.8$ , whereas the standard deviation at  $x_2 = 0.3$   
 410 is significantly higher than at  $x_2 = 0.8$ . Such conflict induces a challenging trade-off  
 411 between performance and robustness.

412 As a consequence, the location of the ground truth Pareto front depends on  
 413 the uncertainty threshold  $\mathbf{p}_v$ : for a high threshold ( $\mathbf{p}_v = 0.5$ ), the Pareto-optimal  
 414 solutions are concentrated around  $x_2 = 0.3$ , prioritising mean performance despite  
 415 higher variability; for a low threshold ( $\mathbf{p}_v = 0.1$ ), the front shifts towards  $x_2 = 0.8$ ,  
 416 favouring lower variability over mean performance.

417 *4.1.2. Searching Strategy*

418 As presented in Section 3 and Eq. (7), the optimisation problem is formulated  
 419 with a hard constraint. However, due to resource limitations, direct evaluation of  
 420 the constraint is infeasible. We anticipate that incorporating the constraint informa-  
 421 tion can still guide the optimisation process towards promising regions of the design  
 422 space. To achieve this, the variance-penalised scalarisation algorithm is proposed  
 423 that transforms the hard constraint into a soft form, enabling control over the influ-  
 424 ence of uncertainty through the parameter  $\alpha$  and thereby steering the search process  
 425 accordingly.

426 Fig. 6 illustrates the distribution of observed points in MDTP3 for different  $\alpha$   
 427 values. When  $\alpha = 0$ , shown in Fig. 6 (a), the search favours regions  $x_2 = 0.3$  with  
 428 low values regardless of variability. As  $\alpha$  increases, the emphasis of the searching  
 429 strategy gradually shifts from  $x_2 = 0.3$  to  $x_2 = 0.8$  where the mean is higher while

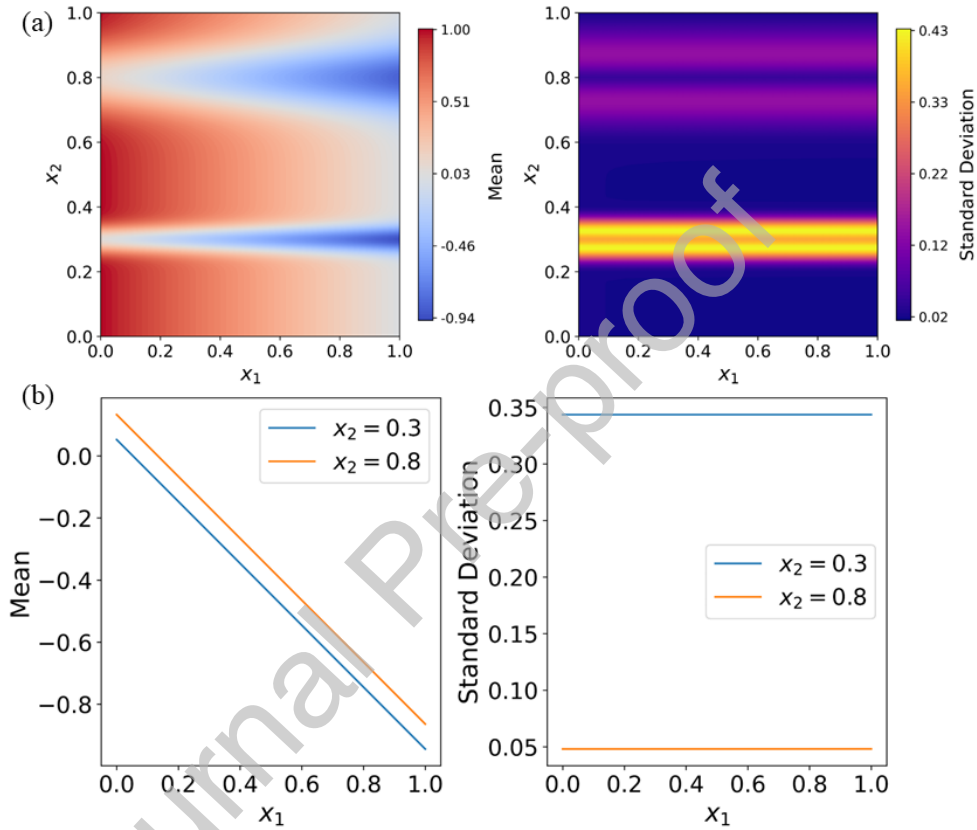


Figure 5: Ground truth performance landscape of MDTP3 objective 2 ( $f_2$ ) under zero-mean Gaussian input noise ( $\sigma_{x_1} = \sigma_{x_2} = 0.02$ ). (a) ( $f_2$ )-mean (left) and -standard deviation (right) over the 2D domain. (b) Cross-sectional statistics at  $x_2 = 0.3$  and  $x_2 = 0.8$ . While  $x_2 = 0.3$  yields a lower mean than  $x_2 = 0.8$ , it also exhibits substantially higher variability, creating a clear trade-off between performance and robustness. This conflict results in distinct Pareto front locations for different uncertainty thresholds: for  $\mathbf{p}_V = 0.5$ , the front is near  $x_2 = 0.3$ ; for  $\mathbf{p}_V = 0.1$ , it shifts towards  $x_2 = 0.8$ .

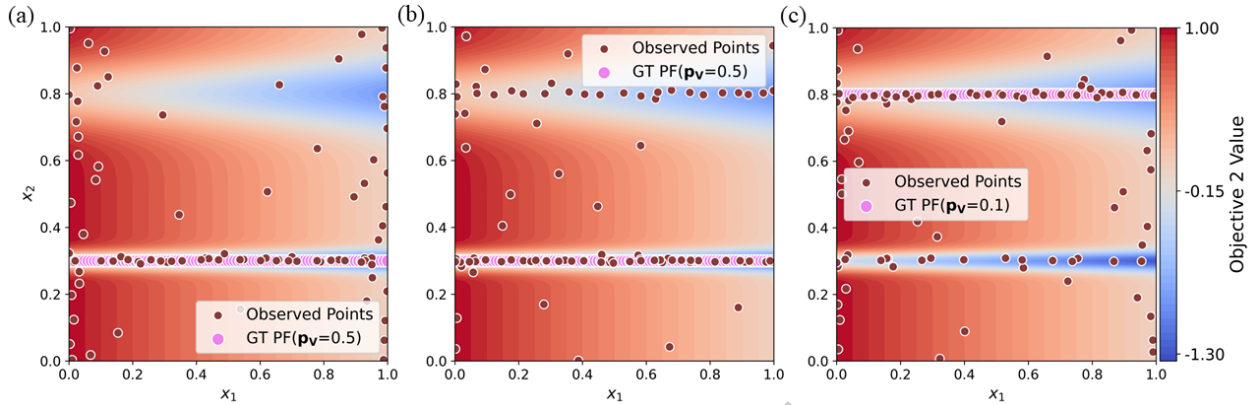


Figure 6: Distribution of observed points (shown in red circles) in the MDTP3 search process under different values of  $\alpha$  in the soft-constrained acquisition function. The pink circles indicate the ground truth Pareto front (GT PF) for the corresponding uncertainty threshold  $\mathbf{p}_V$ . (a)  $\alpha = 0$ . (b)  $\alpha = 0.4$ . (c)  $\alpha = 1$ .

430 uncertainty is lower. By adjusting  $\alpha$  according to the specified constraint threshold,  
 431 the search can be directed to regions of the design space that probably satisfy the  
 432 desired trade-off between performance and robustness. This flexibility allows the  
 433 proposed framework to adapt to different stability requirements without redesigning  
 434 the entire optimisation process, thereby improving efficiency and applicability in  
 435 practical scenarios.

#### 436 4.1.3. Final Performance

437 The VRBO procedure is used to collect data and fit the final surrogates. The  
 438 reported Pareto front is the robust front (RMO), which directly solves the hard-  
 439 constrained problem of Section 3.7 by enforcing the uncertainty threshold  $\mathbf{p}_V$  and  
 440 performing non-dominated sorting on the robust means. For reference, we also post-  
 441 process the final surrogate model to obtain a soft, scalarisation-based front (RMO-  
 442 H): starting from the same GP posteriors used by RMO, we compute the  $H$ -criterion  
 443 given by Eq. 13 for a dense candidate set in the design space and extract the non-  
 444 dominated set in  $\mathbf{H}$ -space. This RMO-H front is shown purely as a comparison  
 445 front to illustrate the effect of scalarising mean and variance; it does not define the  
 446 optimisation target or drive the sampling in our experiments.

447 Additionally, for comparison, an ENSGA-II baseline is also included. ENSGA-  
 448 II adopts a standard NSGA-II procedure operating on a normalised  $[0, 1]^d$  design  
 449 space (population-based evolution, nondominated sorting, crowding-distance selec-  
 450 tion, simulated-binary crossover and polynomial mutation). The only modification  
 451 lies in how each individual is evaluated under input noise. For a given design  $\mathbf{x}$ , we  
 452 draw  $S_{\text{ENSGA}}$  perturbed inputs from the prescribed input-noise distribution centred




Figure 7: Final optimisation results on MDTP3 for different  $\mathbf{p}_V$  values. ENSGA-II (low  $\mathbf{p}_V$ ) and ENSGA-II (high  $\mathbf{p}_V$ ) denote the exhaustive NSGA-II baseline for the respective uncertainty thresholds. RMO: robust Pareto front obtained by applying the final GP to solve the hard-constrained problem. RMO-H: robust Pareto front obtained using the  $H$  criterion with a soft uncertainty constraint. (a)  $\alpha = 0$ ,  $\mathbf{p}_V = 0.5$ , (b)  $\alpha = 0.4$ ,  $\mathbf{p}_V = 0.5$ , and (c)  $\alpha = 1.0$ ,  $\mathbf{p}_V = 0.1$ .

453 at  $\mathbf{x}$ , evaluate the underlying ground-truth benchmark function at each perturbed  
 454 input, and aggregate the resulting objective vectors. The sample means of these  
 455  $S_{\text{ENSGA}}$  realisations are used as the objective values passed to NSGA-II, while the  
 456 corresponding sample standard deviations are used to check whether the design sat-  
 457 isfies the prescribed variability threshold for the given  $\mathbf{p}_V$  level. With a sufficiently  
 458 large exhaustive budget, ENSGA-II provides a high-quality approximation to the  
 459 robust Pareto front. In this paper,  $S_{\text{ENSGA}} = 2000$  to approximate the real mean  
 460 and standard deviation.

461 The results are illustrated in Fig. 7. It can be observed that for  $\alpha = 0$ ,  $\mathbf{p}_V = 0.5$   
 462 (Fig. 7a), both of RMO and ENSGA-II correctly locate the front near  $x_2 \approx 0.3$ .  
 463 RMO-H produces a slight shift due to the soft constraint. While for  $\alpha = 1.0$ ,  $\mathbf{p}_V =$   
 464  $0.1$  (Fig. 7c), all methods correctly locate the front near  $x_2 \approx 0.8$ . In the transition  
 465 setting with  $\alpha = 0.4$  and  $\mathbf{p}_V = 0.5$  (Fig. 7b), RMO-H produces a smooth shift from  
 466 the high- $\mathbf{p}_V$  front towards the low- $\mathbf{p}_V$  front, with the transition initiating from the  
 467 trailing end of the branch, i.e., at larger  $x_1$  values. This is because RMO-H applies  
 468 a soft constraint integrating the mean and standard deviation. Although the entire  
 469 branch exhibits similar uncertainty, the soft formulation of  $H$  makes the importance  
 470 of uncertainty vary with the mean value. In the large- $x_1$  region, where the mean is  
 471 closer to the global minimum, uncertainty will play a more important role, causing  
 472 the shift to begin from this end of the branch.

#### 4.1.4. Evaluation of Performance and Computational Cost

To quantitatively assess the quality of the obtained robust Pareto fronts, three widely used performance indicators are considered. The Generational Distance (GD) [46] measures the average Euclidean distance from each point on the obtained Pareto front to its nearest point on the ground-truth front; smaller GD values indicate better convergence towards the true front. The Inverted Generational Distance (IGD) [46] measures the average distance from each point on the ground-truth front to its nearest point on the obtained front; this metric reflects both convergence and diversity, with smaller values indicating better coverage of the true front. The hypervolume difference ( $\Delta HV$ ) computes the loss in hypervolume between the obtained front and the ground-truth front with respect to a common reference point; a smaller  $\Delta HV$  means the obtained front captures a larger portion of the dominated region defined by the true front.

For each benchmark problem and each variance threshold  $\mathbf{p}_V$ , the ground-truth robust Pareto front is first obtained via a high-resolution grid search over the normalised design space, using the same input-uncertainty level and variance constraint  $\mathbf{p}_V$  as in our experiments. The fronts returned by VRBO and ENSGA-II are then compared against this ground-truth front using GD, IGD and  $\Delta HV$ .

To make the deviations interpretable and comparable across benchmarks, we report GD and IGD as percentages relative to the objective space. Concretely, we first compute GD/IGD in the original objective space and then normalise the resulting Euclidean distances by the size of that space (the Euclidean length of the bounding box defined by the lower and upper bounds of each objective). The normalised values are expressed as percentages, so that a GD or IGD of only a few percent corresponds to average deviations that are small relative to the full objective range. For  $\Delta HV$ , we additionally report the absolute hypervolume of the ground-truth front,  $HV_{GT}$ , which allows the relative impact of  $\Delta HV$  to be judged by comparing it directly to the total attainable hypervolume in each benchmark.

Table 1 summarises, for one representative benchmark, the quantitative results of VRBO and ENSGA-II under two variance thresholds ( $\mathbf{p}_V = 0.5$  and  $\mathbf{p}_V = 0.1$ ), together with the associated computational cost measured as the total number of expensive evaluations of the underlying function. In our sequential surrogate-based setting, these black-box evaluations dominate the overall computational effort, whereas updating and querying the surrogate is comparatively cheap. The number of function evaluations is therefore the most natural and comparable proxy for cost across all methods. The same evaluation protocol (high-resolution grid for ground truth, normalised GD/IGD, and  $\Delta HV$  with  $HV_{GT}$ ) is applied consistently across all benchmark problems considered in this work.

From Table 1, we observe that RMO, evaluated on the VRBO-trained surrogate, attains GD and IGD values on the order of only a few percent of the objective

Table 1: Quantitative comparison between RMO and ENSGA-II in terms of  $\Delta\text{HV}$ , GD, IGD, and computational cost for MDTP3. Lower values of  $\Delta\text{HV}$ , GD, and IGD indicate better performance.

Method	GD(%)	IGD(%)	$\Delta\text{HV}$	$\text{HV}_{\text{GT}}$	Cost (evaluations)
RMO $\alpha = 0$ [ $\mathbf{p}_V = 0.5$ ]	5.7	1.0	0.0151	1.14	90
RMO $\alpha = 0.4$ [ $\mathbf{p}_V = 0.5$ ]	3.4	0.6	0.0151	1.14	85
ENSGA-II [ $\mathbf{p}_V = 0.5$ ]	0.4	0.8	0.0094	1.14	$144 \times 2000$
RMO $\alpha = 1.0$ [ $\mathbf{p}_V = 0.1$ ]	1.4	0.01	0.0105	1.05	73
ENSGA-II [ $\mathbf{p}_V = 0.1$ ]	0.3	0.9	0.0087	1.05	$300 \times 2000$

513 space, together with small  $\Delta\text{HV}$  values relative to the corresponding ground-truth  
 514 hypervolumes  $\text{HV}_{\text{GT}}$  for both variance thresholds  $\mathbf{p}_V = 0.5$  and  $\mathbf{p}_V = 0.1$ . While  
 515 ENSGA-II is included as a reference optimisation baseline, the results indicate that  
 516 RMO is capable of producing robust Pareto fronts with comparable solution quality  
 517 under the considered uncertainty settings.

518 This observation is further supported by the visual comparisons in Fig. 7, where  
 519 the Pareto fronts obtained by RMO are almost indistinguishable from the ground-  
 520 truth fronts obtained via exhaustive search. This efficiency stems from VRBO’s bi-  
 521 loop design and variance-penalised scalarisation (via  $\alpha$ ), which decouple surrogate  
 522 refinement from uncertainty control and steer sampling to stability-feasible regions.

#### 523 4.1.5. Bi-loop Effect

524 All results reported in the previous subsection (e.g., Fig. 7 and Table 1) were ob-  
 525 tained with the full bi-loop VRBO framework. To further examine the contribution  
 526 of the proposed bi-loop strategy, we compare the optimisation results obtained under  
 527 the same  $\alpha$  setting but without the bi-loop. Instead of integrating the traditional  
 528 acquisition function and  $\mathcal{H}$ , the baseline here directly applies the  $\mathcal{H}$  criterion as the  
 529 sole acquisition function throughout the optimisation.

530 The resulting mean and standard-deviation landscapes are shown in Fig. 8. Com-  
 531 pared with the bi-loop results, the single-loop variant exhibits limited exploration  
 532 and a Pareto set that stalls in local minima. This behaviour arises because in the  
 533 RGP model, both the predictive model uncertainty and the propagated input un-  
 534 certainty are intertwined; when they are not decoupled through the bi-loop, the  
 535 optimiser may over-exploit regions with underestimated uncertainty, leading to sub-  
 536 optimal coverage of the true Pareto front.

#### 537 4.2. Additional Benchmarks: VLMOP2 and MDTP2

538 To further examine the generality of the proposed framework, we evaluate it on  
 539 two additional problems: VLMOP2 and MDTP2. These benchmarks represent dis-  
 540 tinct optimisation landscapes, enabling an assessment of performance across both

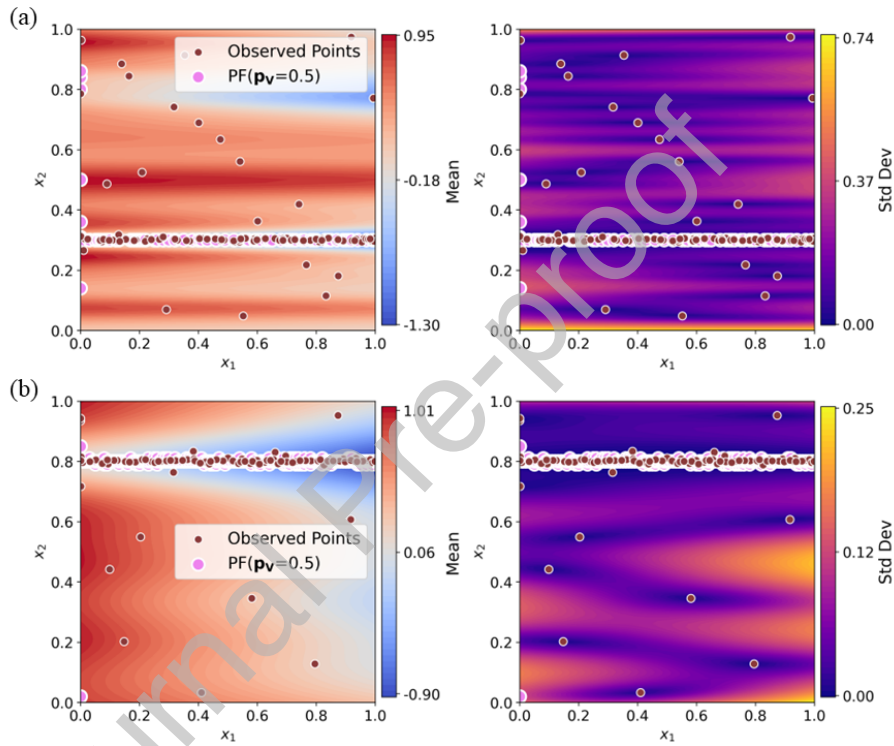


Figure 8: Optimisation results for MDTP3 without the proposed bi-loop strategy, using only the  $H$  criterion as the acquisition function. The lack of alternating updates between model training and acquisition search leads to premature convergence and entrapment in local minima, due to the coupling of model and input uncertainties in RGP. Both examples correspond to  $\alpha = 0$ , illustrating the consistency of this failure mode.

541 smooth and highly nonlinear multi-objective problems. Complete problem formula-  
 542 tions and experimental settings are provided in [Appendix A](#).

#### 543 4.2.1. VLMOP2

544 The VLMOP2 problem is a smooth bi-objective benchmark with a convex and  
 545 continuous Pareto front, commonly used to assess convergence accuracy in noise-free  
 546 multi-objective optimisation. Fig. 9 compares the RMO obtained by the proposed  
 547 VRBO framework and the exhaustive ENSGA-II baseline. The contour maps depict  
 548 the objective function landscapes, with the left panel showing Objective 1 and the  
 549 right panel showing Objective 2. Both methods successfully locate the true Pareto  
 550 front (GT PF), with the VRBO solutions (RMO) closely aligned along the optimal  
 551 trade-off curve. Despite using far fewer function evaluations, the VRBO method  
 552 achieves accuracy comparable to the exhaustive baseline, confirming its efficiency  
 553 even on smooth, well-behaved problems.

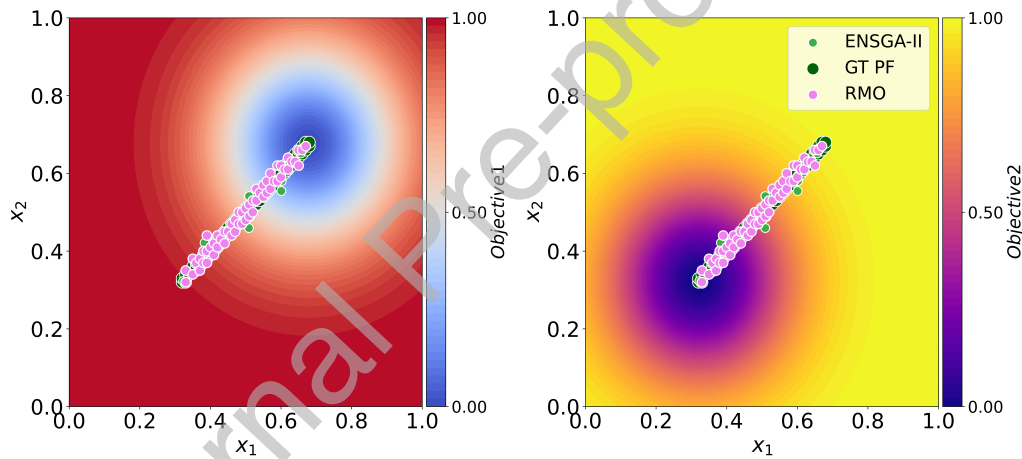


Figure 9: Optimisation results for the VLMOP2 problem: comparison between RMO, ENSGA-II, and the ground-truth Pareto front (GT PF). Contour plots show Objective 1 (left) and Objective 2 (right).

#### 554 4.2.2. MDTP2

555 The MDTP2 problem is a two-objective extension of the Branin-Hoo function,  
 556 featuring strong nonlinearity and multiple local minima in the second objective. This  
 557 structure makes it a more challenging test case, particularly for algorithms that must  
 558 balance global exploration with convergence to the true front. Fig. 10 shows that  
 559 both VRBO (result in RMO) and ENSGA-II identify the ground-truth robust Pareto  
 560 front, with RMO achieving similar accuracy while requiring only a small fraction of  
 561 the evaluations. These results highlight the framework's ability to avoid local optima  
 562 and capture complex front structures under uncertainty.

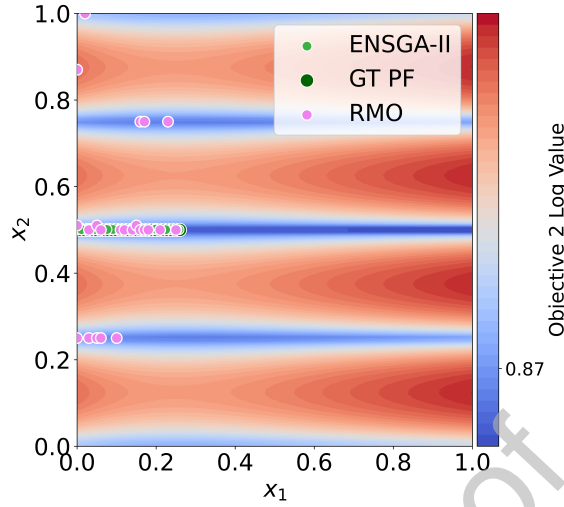


Figure 10: Optimisation results for the MDTP2 problem: comparison between RMO (given by VRBO), ENSGA-II, and the ground-truth Pareto front (GT PF). Contour plots show Objective 2.

56 The final quantitative performance for both problems, including  $\Delta$ HV, GD,  
 563 and IGD, and computational cost, is summarised in Table 2. Lower values of  $\Delta$ HV, GD, and  
 565 IGD indicate better performance.

Table 2: Quantitative comparison between RMO and ENSGA-II in terms of  $\Delta$ HV, GD, IGD, and computational cost for VLMOP2 and MDTP2. Lower values of  $\Delta$ HV, GD, and IGD indicate better performance. Note that a negative  $\Delta$ HV indicates the algorithm outperformed the grid search benchmark, finding denser solutions in certain regions than the discretised baseline.

Benchmark		GD(%)	IGD(%)	$\Delta$ H	HV <sub>G</sub>	Cost
Method		) 0.1	) 0.00	V 0.001	T 0.50	(evaluations)
VLMOP2	RMO [ $\mathbf{p}_v =$ 0.1]	0.6	0.8	0.004	0.50	126 8
	ENSGA-II [ $\mathbf{p}_v =$ 0.1]	6	2	0.133	2.0	2000 9
MDTP	RMO [ $\mathbf{p}_v =$ 2]	35	0	-0.0090	1	72 8
	ENSGA-II [ $\mathbf{p}_v =$ 10]	9	2	2.01		2000

## 566 5. Noisy 567 Observations

568 When optimisation relies on experimental data, different from simulations, mea-  
 569 surement errors, fabrication tolerances, or other perturbations may cause deviations  
 570 between the intended and actual design variables when generating the dataset and  
 571 observation. Therefore, the dataset inherently contains noise. Since experiments  
 572 play an equally important role in engineering design, it is essential to evaluate the  
 proposed framework under these realistic, noisy conditions. Therefore, our method

573 is tested in noisy (experimental) scenarios to comprehensively assess its robustness  
 574 and applicability, as illustrated in Fig. 11.

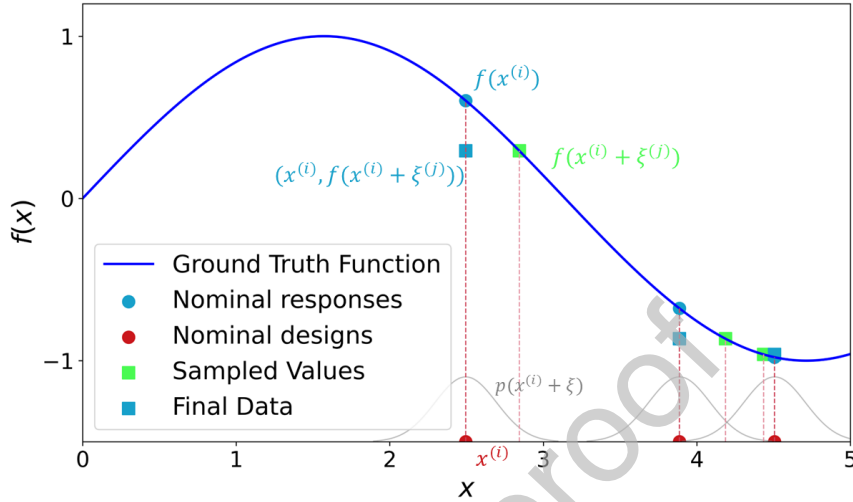


Figure 11: Illustration of input design uncertainty in system response optimisation. Each deterministic proposed design  $x^{(i)}$  (red dot) is in fact perturbed  $x^{(i)} + \xi$  and associated with a distribution, drawn from  $p_\xi(\xi)$  (e.g. shown as grey Gaussian curves). The function evaluation at a particular perturbed location  $f(x^{(i)} + \xi^{(j)})$  (green squares) is therefore wrongly attributed at  $x^{(i)}$  as  $f(x^{(i)})$  (blue squares) which modifies the perceived response and then affects the subsequent optimisation.

### 575 5.1. SinLinForrester

576 SinLinForrester benchmark is utilised to validate the effectiveness under this sce-  
 577 nario.

578 SinLinForrester combines two one-dimensional test objectives, the Sine-Linear  
 579 and Forrester functions, yielding a bi-objective problem in a single-input domain  
 580  $x \in [0, 1]$ . Despite its low dimensionality, it demonstrates conflicting objectives,  
 581 since one has oscillatory behaviour while the other is smoother.

582 The aim for this case is to minimise both  $f_1(x)$  and  $f_2(x)$  defined as

$$\begin{aligned} f_1(x) &= \sin(5\pi x^2) + 0.5x, \\ f_2(x) &= (6x - 2)^2 \sin(12x - 4). \end{aligned} \quad (15)$$

583 Hence, the nominal optimisation problem is  $\min_{0 \leq x \leq 1} [f_1(x), f_2(x)]$ . And Gaussian  
 584 input noise with standard deviation  $\sigma = 0.04$  is introduced to the design parameter.

585 The final results are shown in Fig. 12 and Table 3. In Fig. 12, the solid curves  
 586 depict the ground-truth objectives, the dashed curves show the GP posterior means,  
 587 and the shaded ribbons indicate the input-induced (ground-truth) uncertainty. Ob-  
 588 jective 2 exhibits stronger variability: around  $x \approx 0.80$  it attains lower values but with

589 higher uncertainty, whereas around  $x \approx 0.55$  the trend is reversed. Observed data  
 590 from the optimisation run are overlaid (blue dots for Objective 1, red dots for Objec-  
 591 tive 2), along with vertical ticks illustrating the distribution of sampled inputs. Color  
 592 hue distinguishes the source (blue for MO loop, orange for RMO loop), while color  
 593 saturation encodes the temporal progression, with deeper intensities indicating later  
 594 iterations. Comparing the two scenarios highlights the distinct roles of the two loops  
 595 and  $\alpha$ : the MO loop (indicated by blue ticks) maintains global exploration across  
 596 the entire region, whereas the RMO loop (orange ticks) concentrates on the specific  
 597 area of interest. The location of RMO samples is governed by  $\alpha$ : with a smaller  
 598  $\alpha$  (panel (a),  $\alpha = 0$ ,  $\mathbf{p}_V = 5$ ), the search concentrates near the high-uncertainty  
 599 region around  $x \approx 0.80$  gradually; increasing  $\alpha$  (panel (b),  $\alpha = 0.1$ ,  $\mathbf{p}_V = 2$ ) shifts  
 600 sampling toward safer, lower-uncertainty regions (roughly  $x \approx 0.55$ – $0.65$ ). This be-  
 601 haviour confirms that VRBO flexibly steers exploration–exploitation according to  
 602 the uncertainty preference encoded by  $\alpha$ .

603 Quantitatively, Table 3 compares the final RMO solutions with the exhaustive  
 604 ENSGA-II baseline using  $\Delta HV$ , GD, and IGD against the ground-truth Pareto front.  
 605 For  $\alpha=0$  and  $\mathbf{p}_V=5$ , RMO achieves a small GD because the search concentrates on  
 606 the low-mean yet high-uncertainty region around  $x \approx 0.80$ ; the GP is well trained  
 607 there, so the recovered front lies close to that portion of the ground-truth front.  
 608 However, this focus leaves the branch near  $x \approx 0.55$  sparsely explored, leading to  
 609 under-coverage and hence larger IGD and  $\Delta HV$ . For  $\alpha=0.1$  and  $\mathbf{p}_V=2$ , sampling  
 610 shifts toward safer, lower-uncertainty regions and coverage improves markedly rel-  
 611 ative to  $\alpha=0$  (smaller IGD/ $\Delta HV$ ). Overall, these results highlight how the choice  
 612 of robustness parameter  $\alpha$  and variance threshold  $\mathbf{p}_V$  influences the exploration–  
 613 exploitation balance and the resulting approximation quality of the robust Pareto  
 614 front.

Table 3: Quantitative comparison between RMO and ENSGA-II in terms of  $\Delta HV$ , GD, IGD, and computational cost for SinLinForrester. Lower values of  $\Delta HV$ , GD, and IGD indicate better performance.

Method	GD(%)	IGD(%)	$\Delta HV$	$HV_{GT}$	Cost (evaluations)
RMO $\alpha = 0$ [ $\mathbf{p}_V = 5$ ]	0.34	2.41	0.2366	5.39	45
ENSGA-II [ $\mathbf{p}_V = 5$ ]	1.79	1.11	-0.0367	5.39	$78 \times 2000$
RMO $\alpha = 0.1$ [ $\mathbf{p}_V = 2$ ]	3.08	1.18	0.2749	4.39	45
ENSGA-II [ $\mathbf{p}_V = 2$ ]	1.2	0.61	0.0087	4.39	$108 \times 2000$

615 The fact that the  $\alpha = 0$  run in Fig. 12 does not recover the solution near  $x \approx 0.55$   
 616 reflects the effect of input-uncertainty-corrupted observations on a surrogate-based  
 617 optimiser, rather than a fundamental failure of global exploration; further runs and  
 618 discussion are reported in the Supplementary Material.

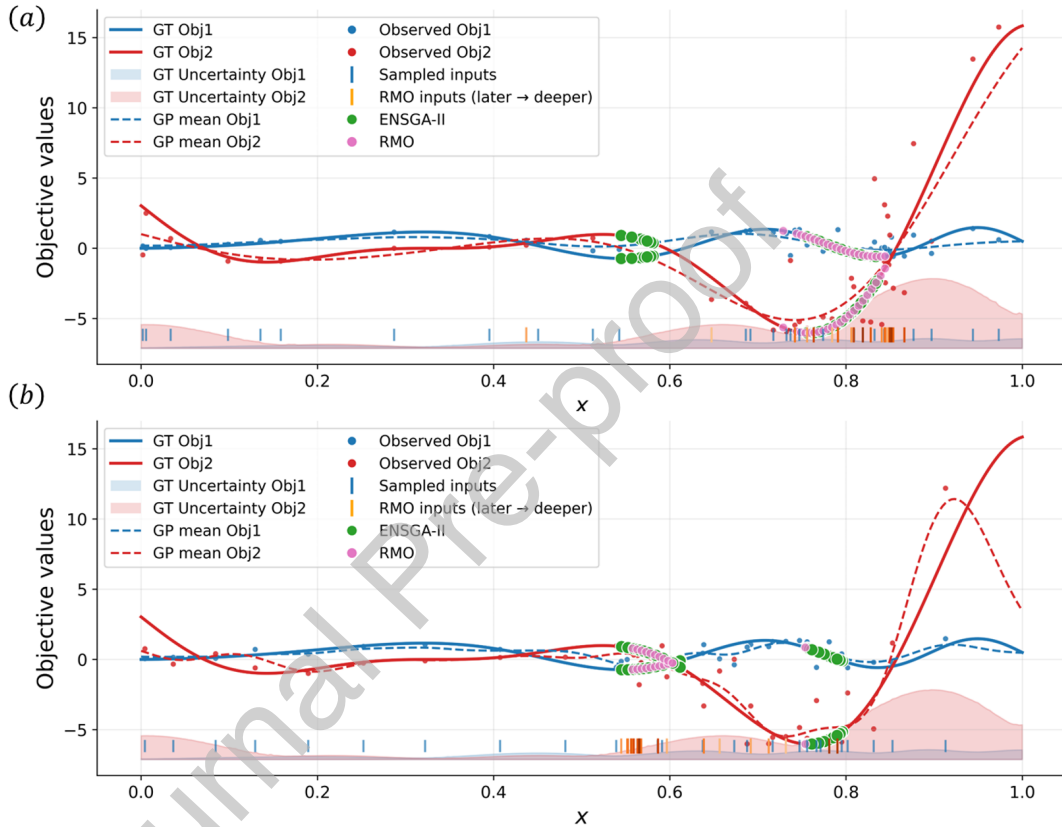


Figure 12: One-dimensional, two-objective example under input noise. Solid curves are the ground-truth objectives; dashed curves are GP posterior means; shaded bands show the ground-truth uncertainty (standard deviation). Dots are observed data collected at nominal  $x$  with noisy outputs (so they do not lie exactly on the ground-truth curves). Short vertical ticks on the  $x$ -axis mark sampled input locations: blue ticks are proposed by the MO loop and orange ticks by the RMO loop (deeper color means later iterations). Green and pink markers denote ENSGA-II and RMO solutions, respectively. (a)  $\alpha = 0$ ,  $\mathbf{pV} = 5$ , (b)  $\alpha = 0.1$ ,  $\mathbf{pV} = 1$ .

## 6. High-Dimension Benchmark case – Fonseca–Fleming

The last test problem is a moderately high-dimensional benchmark based on the Fonseca–Fleming (FON) formulation. It is a smooth bi-objective problem that is frequently used in multi-objective optimisation to assess convergence and diversity in spaces of dimension larger than two. In contrast to the previous two-dimensional benchmarks, this case allows us to investigate how the proposed VRBO framework scales when both the decision space and the input-uncertainty space become higher dimensional.

For a design vector  $\mathbf{x} = (x_1, \dots, x_M)$ , the FON objectives are defined as

$$f_1(\mathbf{x}) = 1 - \exp\left(-\sum_{i=1}^M (x_i - 1/\sqrt{M})^2\right), \quad f_2(\mathbf{x}) = 1 - \exp\left(-\sum_{i=1}^M (x_i + 1/\sqrt{M})^2\right), \quad (16)$$

and the optimisation task is to minimise both objectives

$$\min_{\mathbf{x} \in \mathcal{X}} [f_1(\mathbf{x}), f_2(\mathbf{x})]. \quad (17)$$

In the original definition, the physical design domain is the hypercube  $\mathbf{x} \in [-2, 2]^M$ .

To test robustness under input uncertainty, independent Gaussian perturbations  $\mathcal{N}(\mathbf{0}, \sigma^2)$  are applied to all decision variables in the normalised coordinates, where  $\sigma = 0.02$ . In this study, we consider  $M = 6$  decision variables, a setting that is representative of typical engineering design problems [47, 48].

### 6.1. Ground Truth Analysis

Under this setting, the ground-truth robust Pareto fronts under different uncertainty constraints  $\mathbf{p}_V$  are shown in Fig. 13. To construct this reference, we evaluate the FON on a dense  $21^6$  grid in the normalised design space and, for each design, perform 100 Monte Carlo samples of the input noise. For the chosen noise level, the resulting standard deviation of each objective lies roughly in  $[0, 0.08]$ . Consequently, the looser threshold  $\mathbf{p}_V = 0.10$  (blue dashed curve) is effectively unconstrained and reproduces the nominal Pareto front, whereas the stricter threshold  $\mathbf{p}_V = 0.05$  (orange solid curve) removes high-variance designs and yields a shorter, more conservative front with noticeably higher  $f_2$  for the same  $f_1$  in the middle region. This illustrates how tighter variability constraints shrink the feasible robust Pareto set and push the trade-off towards safer but less optimal designs.

### 6.2. Performance

As presented in Fig. 14 and summarised in Table 4, the behaviour of VRBO under different robustness settings can be clearly observed. When  $\alpha = 0$  and  $\mathbf{p}_V = 1$ , the

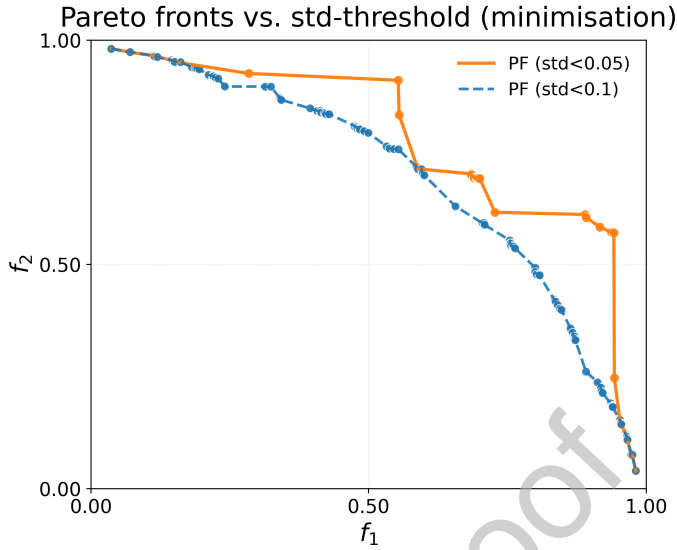


Figure 13: Ground-truth robust Pareto fronts of the FON benchmark under two uncertainty thresholds  $\mathbf{p}_V \in \{0.05, 0.10\}$ . When  $\mathbf{p}_V = 0.1$ , the constraint is not activated. The fronts are obtained by exhaustive evaluation on a  $21^6$  grid with 100 Monte Carlo samples per design.

650 RMO front aligns closely with the GT Pareto front as shown in Fig. 14 (a), with GD  
 651 and IGD both around 0.03, and a small  $\Delta HV$  indicating negligible hypervolume loss.  
 652 However, under the stricter uncertainty requirement  $\mathbf{p}_V = 0.05$ , the approximation  
 653 quality deteriorates noticeably when  $\alpha = 0$  in Fig. 14 (b). Increasing the  $\alpha$  to  $\alpha = 5$   
 654 significantly improves the fit under the same  $\mathbf{p}_V = 0.05$  in Fig. 14 (c), as reflected  
 655 by lower GD, IGD, and  $\Delta HV$ , showing that a larger  $\alpha$  helps recover a robust Pareto  
 656 front closer to the GT reference when tighter variability control is desired.

Table 4: Quantitative comparison between RMO fronts and GT Pareto fronts on the FON problem.

Case	GD(%)	IGD(%)	$\Delta HV$	$HV_{GT}$
$\alpha = 0, \mathbf{p}_V = 1$	3.4	3.4	0.043	0.492
$\alpha = 0, \mathbf{p}_V = 0.05$	7.5	4.6	0.066	0.413
$\alpha = 5, \mathbf{p}_V = 0.05$	4.3	2.5	0.011	0.413

## 657 7. Conclusion

658 This work introduced VRBO, a sample-efficient framework for multi-objective  
 659 optimisation under input uncertainty that unifies robustness and stability within a  
 660 single workflow. The core of VRBO is a RGP surrogate that propagates input per-  
 661 turbations to produce Bayes' risk (means and variances) for each objective. On top of




Figure 14: Comparison between the robust Pareto fronts obtained by VRBO (RMO) and the Monte Carlo ground-truth fronts on the 6D Fonseca–Fleming problem under different robustness parameters. (a)  $\alpha = 0$ ,  $\mathbf{p}_V = 1$ ; (b)  $\alpha = 0$ ,  $\mathbf{p}_V = 0.05$ ; (c)  $\alpha = 5$ ,  $\mathbf{p}_V = 0.05$ .

662 this surrogate, we proposed a variance-penalised scalarisation that balances expected  
 663 performance with uncertainty through a tunable parameter  $\alpha$ , and a bi-loop strat-  
 664 egy that decouples model (epistemic) from objective (aleatory) uncertainty. Stability  
 665 enters the pipeline twice: softly during search via  $\alpha$ , and as a hard uncertainty filter  
 666 using the user tolerance  $\mathbf{p}_V$  when reporting the robust Pareto set.

667 Comprehensive tests of the proposed framework in both noise-free (simulation)  
 668 and noisy (experimental) regimes demonstrate that VRBO reliably locates appropri-  
 669 ate robust Pareto fronts as  $\mathbf{p}_V$  varies. On MDTP3, VRBO reproduces the expected  
 670 shift of the front between high-variance and low-variance bands, and the RMO-H  
 671 variant explains the gradual transition induced by the soft weighting. Across VL-  
 672 MOP2, MDTP2, SinLinForrester, and a six-dimensional Fonseca–Fleming bench-  
 673 mark, the recovered Pareto sets match ground truth to within small GD/IGD/ $\Delta$ HV  
 674 gaps. These results confirm the ability of the proposed framework to accurately  
 675 characterise robust Pareto structures across a range of problem settings, uncertainty  
 676 levels, and dimensionalities.

677 The scalarisation parameter  $\alpha$ , although theoretically related to the uncertainty  
 678 tolerance, still lacks direct interpretability and intuitive selection guidelines. Within  
 679 the current VRBO framework, certain designs (e.g., the variance filter based on  $\mathbf{p}_V$ )  
 680 help to reduce the practical sensitivity to  $\alpha$ , but a clearer connection between  $\mathbf{p}_V$   
 681 and  $\alpha$  remains desirable. Future work will therefore focus on establishing a more  
 682 explicit link between robustness tolerance and scalarisation, potentially through a  
 683 normalisation scheme that treats mean performance and uncertainty on a compara-  
 684 ble scale, and on increasing the adaptivity of the acquisition so that  $\alpha$  can evolve  
 685 automatically with the search process as a function of  $\mathbf{p}_V$ . In parallel, we also pro-  
 686 vide in the [Appendix B](#) a supplementary split-acquisition variant that does not

687 rely on  $\alpha$  and instead incorporates  $\mathbf{p}_v$  directly into the optimisation objectives, of-  
 688 fering an alternative way to encode robustness within the algorithm. Beyond these  
 689 methodological developments, applying the framework to an industrial battery heat-  
 690 exchanger design problem is the next step toward demonstrating the practical utility  
 691 of VRBO in real engineering settings.

## 692 Data and code availability

693 The code and example structure of models used in this paper is available at  
 694 <https://github.com/HaoZhou-713/rbro-vrbo.git>.

## 695 Declaration of interests

696 The authors declare the following financial interests/personal relationships which  
 697 may be considered as potential competing interests:

698 Emilie Sauret reports financial support was provided by Australian Research  
 699 Council. If there are other authors, they declare that they have no known competing  
 700 financial interests or personal relationships that could have appeared to influence the  
 701 work reported in this paper.

## 702 Acknowledgement

703 This work was supported by the Australian Research Council Discovery Project  
 704 DP230102229 awarded to Emilie Sauret. Hao Zhou gratefully acknowledges the QUT  
 705 for support through a Ph.D. scholarship.

## 706 Abbreviations

707	<b>BO</b>	Bayesian Optimisation
708	<b>CVaR</b>	Conditional Value-at-Risk
709	<b>EHVI</b>	Expected Hypervolume Improvement
710	<b>EI</b>	Expected Improvement
711	<b>ENSGA-II</b>	Exhausted NSGA-II
712	<b>FON</b>	Fonseca-Flemin
713	<b>GD</b>	Generational Distance
714	<b>GP</b>	Gaussian Process
715	<b>IGD</b>	Inverted Generational Distance

716	<b>MOBO</b>	Multi-Objective Bayesian Optimisation
717	<b>MOO</b>	Multi-Objective Optimisation
718	<b>NSGA-II</b>	Non-Dominated Sorting Genetic Algorithm II
719	<b>RBDO</b>	Reliability-based Design Optimisation
720	<b>RBO</b>	Robust Bayesian Optimisation
721	<b>RBRDO</b>	Reliability-based Robust Design Optimisation
722	<b>RDO</b>	Robust-Design Optimisation
723	<b>RGP</b>	Robust Gaussian Process
724	<b>RO</b>	Robust Optimisation
725	<b>UCB</b>	Upper Confidence Boundary
726	<b>VaR</b>	Value-at-Risk
727	<b>VRBO</b>	Variance-constrained Robust Bayesian Optimisation

## 728 Appendix A. Benchmark Definitions: VLMOP2 and MDTP2

729 This appendix specifies the exact formulations, domains, and input-noise settings  
 730 used for the two auxiliary benchmarks. In both cases we conduct optimisation and  
 731 robust evaluation in the normalised coordinate system  $\mathbf{x} = (x_1, x_2) \in [0, 1]^2$ , and  
 732 apply independent Gaussian input perturbations to each component:

$$\xi \sim \mathcal{N}(\mathbf{0}, (0.02)^2 \mathbf{I}_2), \quad \text{i.e., } x_j \leftarrow x_j + \xi_j.$$

733 Unless stated otherwise in the main text, robust means and variances (or standard  
 734 deviations) are estimated by Monte Carlo under this noise model; feasibility is then  
 735 enforced via the uncertainty tolerance used in each experiment before extracting the  
 736 non-dominated set on robust means.

### 737 Appendix A.1. VLMOP2

738 We use the standard VLMOP2 bi-objective problem defined on the physical co-  
 739 ordinates  $\mathbf{u} = (u_1, u_2) \in [-2, 2]^2$ :

$$f_1(\mathbf{u}) = 1 - \exp\left(-\left(u_1 - \frac{1}{\sqrt{2}}\right)^2 - \left(u_2 - \frac{1}{\sqrt{2}}\right)^2\right),$$

$$f_2(\mathbf{u}) = 1 - \exp\left(-\left(u_1 + \frac{1}{\sqrt{2}}\right)^2 - \left(u_2 + \frac{1}{\sqrt{2}}\right)^2\right).$$

740 Optimisation is carried out in the normalised domain  $\mathbf{x} \in [0, 1]^2$ . For function  
 741 evaluation, we map  $\mathbf{x}$  back to the physical domain  $\mathbf{u}$ , then compute  $f_1, f_2$ . Input  
 742 noise is applied in the normalised coordinates (STD = 0.02 per dimension) prior to  
 743 evaluation. The uncertainty tolerance used for VLMOP2 is as stated in the main  
 744 text (see Table 2).

#### 745 *Appendix A.2. MDTP2*

746 The MDTP2 bi-objective variant used in this work is

$$\begin{aligned} \mathbf{u} &= (u_1, u_2), & u_1 &\in [0, 1], & u_2 &\in [-1, 1], \\ f_1(\mathbf{u}) &= u_1, \\ f_2(\mathbf{u}) &= \left(1 - u_1^2\right) + \left(10 + u_2^2 - 10 \cos(4\pi u_2)\right) \left(\frac{1}{0.2+u_1} + 10u_1^2\right). \end{aligned}$$

747 As with VLMOP2, optimisation and robust evaluation are performed in the nor-  
 748 malised space  $\mathbf{x}$ ; independent Gaussian input noise with STD 0.02 is applied to each  
 749 normalised coordinate. The uncertainty tolerance and evaluation details match those  
 750 reported in the main text for this case (e.g., Table 2).

#### 751 **Appendix B. Alternative Split Formulation for Parameter-Free Optimi-** 752 **sation**

753 To address the challenge of selecting the scalarisation parameter  $\alpha$  in Eq. 13 and  
 754 to provide a parameter-free alternative, this appendix presents a split formulation  
 755 where the mean performance and the uncertainty are treated as two explicit, separate  
 756 objectives. Recall that the variance-penalised scalarisation for the  $i$ -th objective is  
 757 originally defined as  $\mathcal{H}_i(x) = \text{CVaR}_{\theta_i}(x) + \alpha \ln(1 + \sigma_i(x))$ . In the split formulation,  
 758 this scalarised acquisition is decoupled into two distinct objectives to be minimised  
 759 simultaneously: the mean performance, denoted as  $f_i^{(1)}(x) = \text{CVaR}_{\theta_i}(x)$ , and the  
 760 normalised uncertainty, denoted as  $f_i^{(2)}(x)$ .

761 To ensure the uncertainty objective is commensurate with the user-specified vari-  
 762 ance tolerance  $p_{V,i}$ , a normalisation factor is introduced. The uncertainty objective  
 763 is thus defined as:

$$f_i^{(2)}(x) = \frac{\ln(1 + \sigma_i(x))}{\ln(1 + p_{V,i})}. \quad (\text{B.1})$$

764 This mapping scales the uncertainty term relative to the allowable variability budget,  
 765 thereby removing the need for the manual trade-off weight  $\alpha$ . Consequently, for an  
 766 original problem with  $N$  objectives, the optimisation transforms into a  $2N$ -objective  
 767 problem:

$$\min_{x \in \mathcal{X}} \left( \text{CVaR}_{\theta_1}(x), \dots, \text{CVaR}_{\theta_N}(x), f_1^{(2)}(x), \dots, f_N^{(2)}(x) \right). \quad (\text{B.2})$$

768 The split acquisition formulation was evaluated on two benchmark problems:  
 769 MDTP3 and the noisy SinLinForrester case. The framework demonstrates the antici-  
 770 pated behavior when the variability tolerance is stringent. Specifically, with  $\mathbf{p}_V = 0.1$   
 771 on MDTP3 (Fig. B.15(a)) and  $\mathbf{p}_V = 1$  on SinLinForrester (Fig. B.16(a)), the sam-  
 772 pler concentrates on regions satisfying the prescribed variability constraints while  
 773 simultaneously exploiting high-quality areas of the design space.

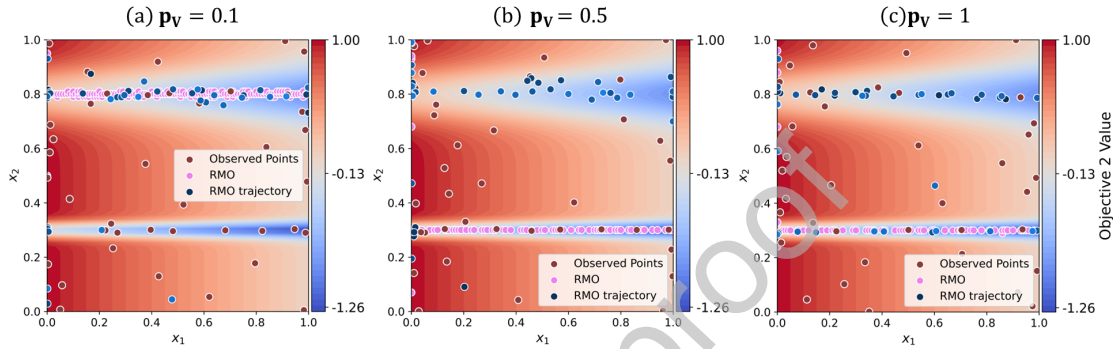


Figure B.15: Distribution of sampled points in the MDTP3 search under different values of  $\mathbf{p}_V$ . Red circles denote points proposed by the MO loop, blue circles (darker shades indicate later iterations) denote points proposed by the RMO loop, and pink circles indicate the robust Pareto front corresponding to the specific uncertainty threshold  $\mathbf{p}_V$ .

774 However, moderate increases in  $\mathbf{p}_V$  result in less discernible changes in sampling  
 775 behavior. As observed with  $\mathbf{p}_V = 0.5$  on MDTP3 (Fig. B.15(b)) and  $\mathbf{p}_V = 5$   
 776 on SinLinForrester (Fig. B.16(b)), the sampling locations remain similar to those  
 777 obtained with tighter tolerances. Significant shifts in sampling behaviour—specifically,  
 778 the allocation of points into higher-uncertainty bands—are observed only when  $\mathbf{p}_V$   
 779 is substantially increased (e.g.,  $\mathbf{p}_V = 1$  on MDTP3 as shown in Fig. B.15(c), and  
 780  $\mathbf{p}_V = 10$  on SinLinForrester as shown in Fig. B.16(c)).

781 This phenomenon is attributed to the problem-dependent uncertainty landscape.  
 782 For instance, in the SinLinForrester case, the maximum ground-truth standard de-  
 783 viation is approximately 5 (see Fig. B.17). Consequently, setting  $\mathbf{p}_V = 5$  results in  
 784 the normalisation of the term

$$\text{unc}(x) = \frac{\ln(1 + \sigma(x))}{\ln(1 + \mathbf{p}_V)}$$

785 to the interval  $[0, 1]$ . This functions similarly to a standard min-max normalisation,  
 786 causing the acquisition function to retain a preference for lower-uncertainty regions.  
 787 Conversely, increasing  $\mathbf{p}_V$  to 10 results in a denominator that exceeds the maximum  
 788 realisable data value. This effectively compresses the typical uncertainty scores to-  
 789 ward zero, thereby diminishing the relative influence of the uncertainty objective.

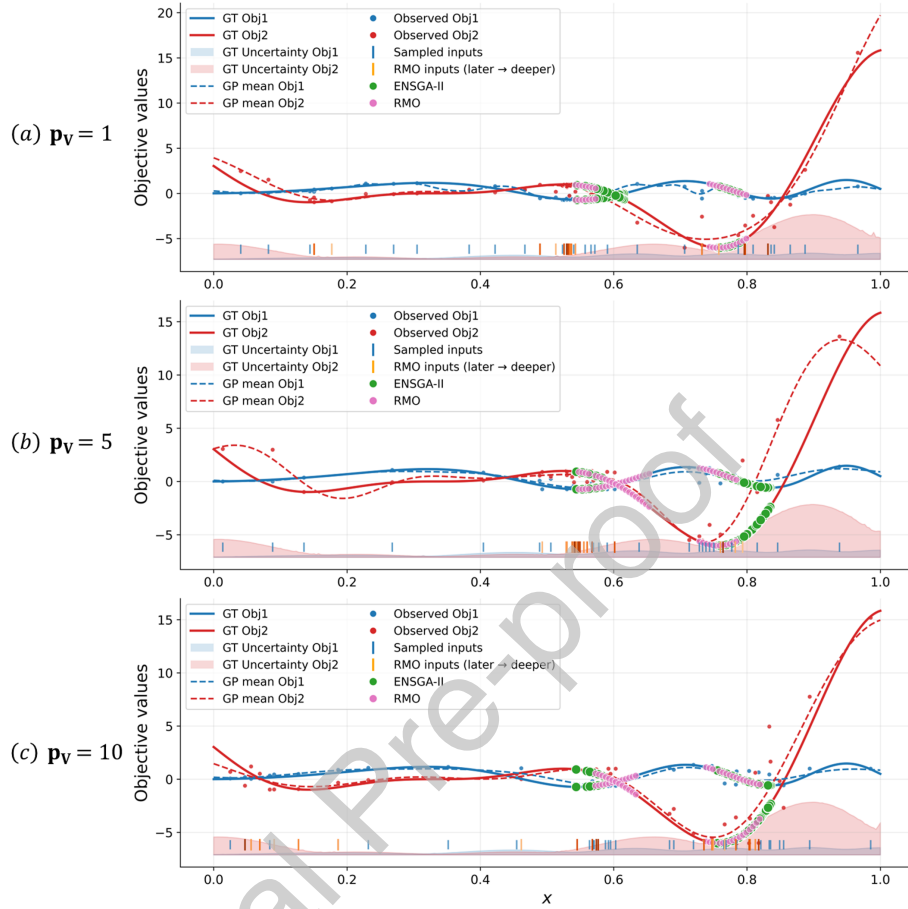


Figure B.16: One-dimensional, two-objective example under input noise. Solid curves represent the ground-truth objectives; dashed curves indicate GP posterior means; shaded bands show the ground-truth uncertainty (standard deviation). Dots represent observed data collected at nominal  $x$  with noisy outputs. Short vertical ticks on the  $x$ -axis mark sampled input locations: blue ticks correspond to the MO loop and orange ticks to the RMO loop (deeper color intensity indicates later iterations). Green and pink markers denote ENSGA-II and RMO solutions, respectively. (a)  $\mathbf{p}_V = 1$ , (b)  $\mathbf{p}_V = 5$ , (c)  $\mathbf{p}_V = 10$ .

790 As a result, the subsequent EHVI calculation penalises moderate uncertainty less  
 791 severely, facilitating selection in higher-uncertainty regions.

792 In summary, while the split formulation is feasible and eliminates the need for  
 793 manual tuning of  $\alpha$ , effective practical implementation suggests setting  $\mathbf{p}_V$  slightly  
 794 above the expected system variability (approximately 1–2 $\times$ ) to ensure the uncer-  
 795 tainty objective remains sufficiently active.

796 A critical consideration for this approach is the associated computational cost.  
 797 Decomposing each objective into separate CVaR and unc components doubles the

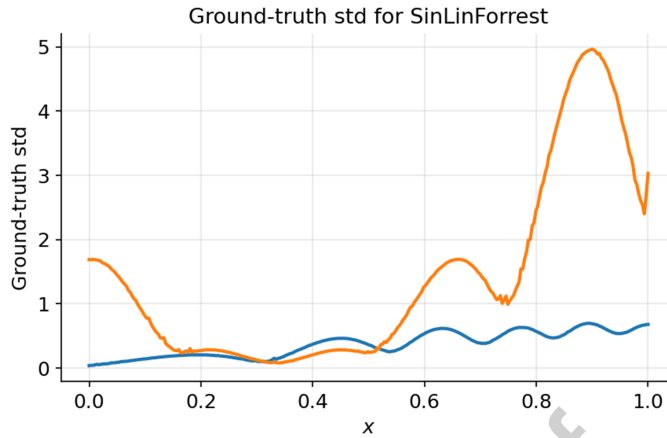


Figure B.17: Ground-truth input-noise profiles for the SinLinForrester benchmark. The orange curve shows the standard deviation of the disturbed objective, which varies with  $x$ , while the blue curve corresponds to the lower-variance objective.

798 dimensionality of the objective space passed to the HV computation (from  $d$  to  $2d$ ).  
 799 Exact HV calculation algorithms are highly sensitive to dimension: complexity typ-  
 800 ically scales from  $O(n \log n)$  or  $O(n^2)$  for  $d = 2, 3$  to polynomial time  $O(n^{f(d)})$  for  
 801  $d \geq 4$ , where the exponent grows with  $d$ . In the conducted experiments, this dimen-  
 802 sionality shift increased the runtime from approximately 5 minutes to 23 minutes.  
 803 This overhead is an inherent trade-off for making the uncertainty trade-off explicit;  
 804 however, it may be mitigated in practice through candidate pre-filtering or the use  
 805 of approximate HV algorithms.

806 **Bibliography**

- 807 [1] J. L. J. Pereira, G. A. Oliver, M. B. Francisco, S. S. Cunha Jr, G. F. Gomes, A review of  
808 multi-objective optimization: methods and algorithms in mechanical engineering problems,  
809 *Archives of Computational Methods in Engineering* 29 (2022) 2285–2308.
- 810 [2] I. Rahimi, A. H. Gandomi, F. Chen, E. Mezura-Montes, A review on constraint handling tech-  
811 niques for population-based algorithms: from single-objective to multi-objective optimization,  
812 *Archives of Computational Methods in Engineering* 30 (2023) 2181–2209.
- 813 [3] B. Y. Qu, P. N. Suganthan, Constrained multi-objective optimization algorithm with an  
814 ensemble of constraint handling methods, *Engineering Optimization* 43 (2011) 403–416.
- 815 [4] K. Deb, Multi-objective optimisation using evolutionary algorithms: an introduction, in:  
816 *Multi-objective evolutionary optimisation for product design and manufacturing*, Springer,  
817 2011, pp. 3–34.
- 818 [5] E. Zitzler, M. Laumanns, L. Thiele, Spea2: Improving the strength pareto evolutionary algo-  
819 rithm, *TIK report* 103 (2001).
- 820 [6] J. Luo, F. Liu, Performance impact of manufacturing tolerances for a turbine blade using  
821 second order sensitivities, in: *Turbo Expo: Power for Land, Sea, and Air*, volume 51029,  
822 American Society of Mechanical Engineers, 2018, p. V02DT46A010.
- 823 [7] M. A. Elfarra, M. N. Önel, G. Yalin, S. Malik, Impact of production tolerances on turbine  
824 performance: A cfd based monte-carlo simulations study, *Alexandria Engineering Journal* 117  
825 (2025) 289–300.
- 826 [8] S. Sudhindra, F. Kargar, A. A. Balandin, Noncured graphene thermal interface materials for  
827 high-power electronics: Minimizing the thermal contact resistance, *Nanomaterials* 11 (2021)  
828 1699.
- 829 [9] J. Ide, A. Schöbel, Robustness for uncertain multi-objective optimization: a survey and analysis  
830 of different concepts, *OR spectrum* 38 (2016) 235–271.
- 831 [10] J. R. Birge, F. Louveaux, *Introduction to stochastic programming*, Springer, 1997.
- 832 [11] S. Mandal, A. Shiuly, D. Sau, A. K. Mondal, K. Sarkar, Study on the use of different machine  
833 learning techniques for prediction of concrete properties from their mixture proportions with  
834 their deterministic and robust optimisation, *AI in Civil Engineering* 3 (2024) 7.
- 835 [12] S. Daulton, S. Cakmak, M. Balandat, M. A. Osborne, E. Zhou, E. Bakshy, Robust multi-  
836 objective bayesian optimization under input noise, in: *International Conference on Machine*  
837 *Learning*, PMLR, 2022, pp. 4831–4866.
- 838 [13] H. Cheng, Z. Zhang, et al., Aerodynamic robustness optimization of aeroengine fan perfor-  
839 mance based on an interpretable dynamic machine learning method, *Reliability Engineering*  
840 *& System Safety* 254 (2025) 110654.
- 841 [14] W. Hu, S. Cheng, J. Yan, J. Cheng, X. Peng, H. Cho, I. Lee, Reliability-based design optimiza-  
842 tion: a state-of-the-art review of its methodologies, applications, and challenges, *Structural*  
843 *and Multidisciplinary Optimization* 67 (2024) 168.

- 844 [15] D. Meng, H. Yang, S. Yang, Y. Zhang, A. M. De Jesus, J. Correia, T. Fazerer-Ferradosa,  
845 W. Macek, R. Branco, S.-P. Zhu, Kriging-assisted hybrid reliability design and optimization  
846 of offshore wind turbine support structure based on a portfolio allocation strategy, *Ocean*  
847 *Engineering* 295 (2024) 116842.
- 848 [16] J. Song, Y. Cui, P. Wei, M. A. Valdebenito, W. Zhang, Constrained bayesian optimization  
849 algorithms for estimating design points in structural reliability analysis, *Reliability Engineering*  
850 *& System Safety* 241 (2024) 109613.
- 851 [17] R. d. S. Motta, S. M. Afonso, An efficient procedure for structural reliability-based robust  
852 design optimization, *Structural and Multidisciplinary Optimization* 54 (2016) 511–530.
- 853 [18] M. Yang, H. Zhang, D. Zhang, X. Han, Q. Li, Time-variant reliability-based robust optimiza-  
854 tion for structures with material degradation, *Computer Methods in Applied Mechanics and*  
855 *Engineering* 432 (2024) 117337.
- 856 [19] S. Sarangi, K. K. Bodla, S. V. Garimella, J. Y. Murthy, Manifold microchannel heat sink  
857 design using optimization under uncertainty, *International Journal of Heat and Mass Transfer*  
858 69 (2014) 92–105.
- 859 [20] D. Fernández-Sánchez, E. C. Garrido-Merchán, D. Hernández-Lobato, Improved max-value  
860 entropy search for multi-objective bayesian optimization with constraints, *arXiv preprint*  
861 *arXiv:2011.01150* (2020).
- 862 [21] D. R. Jones, M. Schonlau, W. J. Welch, Efficient global optimization of expensive black-box  
863 functions, *Journal of Global optimization* 13 (1998) 455–492.
- 864 [22] C. E. Rasmussen, Gaussian processes in machine learning, in: *Summer school on machine*  
865 *learning*, Springer, 2003, pp. 63–71.
- 866 [23] Q. Zhou, P. Jiang, X. Huang, F. Zhang, T. Zhou, A multi-objective robust optimization  
867 approach based on gaussian process model, *Structural and Multidisciplinary Optimization* 57  
868 (2018) 213–233.
- 869 [24] J. Qing, I. Couckuyt, T. Dhaene, A robust multi-objective bayesian optimization framework  
870 considering input uncertainty, *Journal of Global Optimization* 86 (2023) 693–711.
- 871 [25] Z. Abulawi, R. Hu, P. Balaprakash, Y. Liu, Bayesian optimized deep ensemble for uncertainty  
872 quantification of deep neural networks: a system safety case study on sodium fast reactor  
873 thermal stratification modeling, *Reliability Engineering & System Safety* (2025) 111353.
- 874 [26] A. Kudva, W.-T. Tang, J. A. Paulson, Robust bayesian optimization for flexibility analysis of  
875 expensive simulation-based models with rigorous uncertainty bounds, *Computers & Chemical*  
876 *Engineering* 181 (2024) 108515.
- 877 [27] Z. Li, K. Fu, B. Bevans, P. Rao, R. Wang, A. Carrington, A. Riensche, C. Barrett, S. Halliday,  
878 M. Kim, Scalable and robust bayesian transfer learning for in-situ qualification in laser powder  
879 bed fusion additive manufacturing, *Authorea Preprints* (2025).
- 880 [28] R. Oliveira, L. Ott, F. Ramos, Bayesian optimisation under uncertain inputs, in: *The 22nd*  
881 *international conference on artificial intelligence and statistics*, PMLR, 2019, pp. 1177–1184.

- 882 [29] Y. Inatsu, S. Takeno, H. Hanada, K. Iwata, I. Takeuchi, Bounding box-based multi-objective  
883 bayesian optimization of risk measures under input uncertainty, in: International Conference  
884 on Artificial Intelligence and Statistics, PMLR, 2024, pp. 4564–4572.
- 885 [30] J. Goh, M. Sim, Distributionally robust optimization and its tractable approximations, *Op-*  
886 *erations research* 58 (2010) 902–917.
- 887 [31] D. Bertsimas, D. B. Brown, C. Caramanis, Theory and applications of robust optimization,  
888 *SIAM review* 53 (2011) 464–501.
- 889 [32] R. T. Rockafellar, S. Uryasev, et al., Optimization of conditional value-at-risk, *Journal of risk*  
890 2 (2000) 21–42.
- 891 [33] X. Chen, M. Sim, P. Sun, A robust optimization perspective on stochastic programming,  
892 *Operations research* 55 (2007) 1058–1071.
- 893 [34] S. Kang, K. Li, R. Wang, A survey on pareto front learning for multi-objective optimization,  
894 *Journal of Membrane Computing* (2024) 1–7.
- 895 [35] T. M. Deist, M. Grewal, F. J. Dankers, T. Alderliesten, P. A. Bosman, Multi-objective learn-  
896 ing using hv maximization, in: International Conference on Evolutionary Multi-Criterion  
897 Optimization, Springer, 2023, pp. 103–117.
- 898 [36] G. Su, L. Peng, L. Hu, A gaussian process-based dynamic surrogate model for complex  
899 engineering structural reliability analysis, *Structural Safety* 68 (2017) 97–109.
- 900 [37] K. Zhang, A. Guliani, S. Ogrenç-Memik, G. Memik, K. Yoshii, R. Sankaran, P. Beckman,  
901 Machine learning-based temperature prediction for runtime thermal management across system  
902 components, *IEEE Transactions on parallel and distributed systems* 29 (2017) 405–419.
- 903 [38] J. Wilson, F. Hutter, M. Deisenroth, Maximizing acquisition functions for bayesian optimiza-  
904 tion, *Advances in neural information processing systems* 31 (2018).
- 905 [39] W. Gan, Z. Ji, Y. Liang, Acquisition functions in bayesian optimization, in: 2021 2nd  
906 international conference on big data & artificial intelligence & software engineering (ICBASE),  
907 IEEE, 2021, pp. 129–135.
- 908 [40] S. Daulton, M. Balandat, E. Bakshy, Differentiable expected hypervolume improvement for  
909 parallel multi-objective bayesian optimization, *Advances in Neural Information Processing*  
910 *Systems* 33 (2020) 9851–9864.
- 911 [41] K. Yang, M. Emmerich, A. Deutz, T. Bäck, Efficient computation of expected hypervolume  
912 improvement using box decomposition algorithms, *Journal of Global Optimization* 75 (2019)  
913 3–34.
- 914 [42] J. Qing, T. Dhaene, I. Couckuyt, Spectral representation of robustness measures for optimiza-  
915 tion under input uncertainty, in: 38th International Conference on Machine Learning (ICML),  
916 2022.
- 917 [43] L. Yang, J. Lyu, W. Lyu, Z. Chen, Efficient robust bayesian optimization for arbitrary uncer-  
918 tain inputs, *Advances in Neural Information Processing Systems* 36 (2023) 43156–43176.

- 919 [44] Y. Chow, A. Tamar, S. Mannor, M. Pavone, Risk-sensitive and robust decision-making: a cvar  
920 optimization approach, *Advances in neural information processing systems* 28 (2015).
- 921 [45] M. Balandat, B. Karrer, D. Jiang, S. Daulton, B. Letham, A. G. Wilson, E. Bakshy, Botorch:  
922 A framework for efficient monte-carlo bayesian optimization, *Advances in neural information*  
923 *processing systems* 33 (2020) 21524–21538.
- 924 [46] H. Ishibuchi, H. Masuda, Y. Tanigaki, Y. Nojima, Modified distance calculation in generational  
925 distance and inverted generational distance, in: *International conference on evolutionary multi-*  
926 *criterion optimization*, Springer, 2015, pp. 110–125.
- 927 [47] A. Verma, T. Saikia, P. Saikia, D. Rakshit, C. E. Ugalde-Loo, Thermal performance analysis  
928 and experimental verification of lithium-ion batteries for electric vehicle applications through  
929 optimized inclined mini-channels, *Applied Energy* 335 (2023) 120743.
- 930 [48] W. Li, M. Xiao, X. Peng, A. Garg, L. Gao, A surrogate thermal modeling and parametric  
931 optimization of battery pack with air cooling for evs, *Applied Thermal Engineering* 147 (2019)  
932 90–100.