MEAN–VARIANCE PARETO LEARNING WITH REPLICATE-AWARE BAYESIAN OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Stochastic simulators often exhibit a trade-off between average performance and run-to-run variability. When the preferred balance is not known in advance, the goal is not one weighted compromise but the set of efficient designs spanning this trade-off, namely the mean–variance Pareto frontier. Learning this frontier under a limited simulation budget creates an additional difficulty: at each design, the mean and variance are estimated from the same finite collection of simulator runs. Their estimation errors therefore depend on the number of replications and can be correlated, so a sequential algorithm must decide both where to sample and how much simulation effort to allocate there. We develop a replicate-aware Bayesian optimization framework for this joint learning and allocation problem. We first derive the exact covariance of the sample mean and sample variance for any finite replication count and establish their joint asymptotic distribution. We then embed this sampling structure in a hierarchical multi-output Gaussian process that separates finite-sample estimation dependence from dependence between the underlying mean and variance surfaces. Based on this posterior, hierarchical expected quantile hypervolume improvement (HEQHVI) jointly selects a design and replication count by maximizing the expected hypervolume gain of a posterior-quantile Pareto front per simulator run. This creates an explicit trade-off between exploring new designs and reducing uncertainty at previously evaluated ones. For parallel evaluation, we extend HEQHVI using a minimum-energy design (MED) criterion that retains the best serial candidate while spreading the remaining batch members across promising but distinct regions. Synthetic benchmarks and an F119 aero-engine study illustrate how the framework learns mean–variance trade-offs while adapting replication effort.

1 INTRODUCTION

Many important optimization problems in science, engineering, and machine learning involve expensive black-box evaluations, such as high-fidelity simulations and costly learning pipelines. Bayesian optimization (BO) provides an efficient framework for such problems by using a probabilistic surrogate to guide evaluations when exhaustive search is infeasible. A further challenge arises when the black-box output is stochastic: repeated evaluations of the same input can produce different outcomes because of simulation randomness, training randomness, environmental variation, or operational uncertainty (Qi et al., 2023; Zhou et al., 2023; Lejeune & Ma, 2025). In this setting, an input is characterized not by a single performance value but by a distribution of possible outcomes. Optimizing only the mean can therefore be insufficient, since two designs with similar average performance may exhibit very different run-to-run variability, while a design with a better mean may also be substantially less stable. This motivates considering the mean and variance jointly, so that optimization can account for both average performance and variability.

Existing methods for stochastic optimization address output variability in several ways. Stochastic BO and simulation-optimization methods model noisy or heteroscedastic responses and may allocate repeated evaluations to improve estimation accuracy (Wang et al., 2025; Ankenman et al., 2010; Binois et al., 2018; Jalali et al., 2017; Binois et al., 2019; Binois & Larson, 2026). Risk-aware BO further incorporates variability into the optimization objective, for example through quantiles, risk measures, or other scalar criteria (Picheny et al., 2013; Makarova et al., 2021; Griffiths et al., 2022; Cakmak et al., 2020; Tay et al., 2024). Mean–variance optimization has also been studied

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

directly (Iwazaki et al., 2021), while adaptive replication has been considered in batch experimental design (Dai et al., 2023). A scalar objective is appropriate when the desired trade-off between average performance and variability is known in advance. When that preference is unknown, however, committing to one particular trade-off can discard designs that would be attractive under other preferences. Learning the mean–variance Pareto frontier instead preserves the set of efficient trade-off solutions without selecting a single compromise beforehand.

Multi-objective BO provides a natural framework for learning such a frontier. Methods such as ParEGO and qNParEGO explore multiple objectives through random scalarizations, while hypervolume-based methods directly value the improvement of the current Pareto set (Knowles, 2006; Daulton et al., 2020; 2021). Quantile-based improvement has also been extended to multi-objective optimization (Semochkina et al., 2025). These methods, however, do not by themselves address the observation structure that arises when the two objectives are the mean and variance estimated from the same finite set of simulator runs. Because the same replications are used to compute both quantities, their estimation errors can be correlated, and their uncertainty changes with the number of replications. Moreover, the replication count is itself a decision: additional runs at an existing design improve estimation precision but consume simulation budget that could instead be used to explore a new design. Thus, mean–variance Pareto learning under stochastic simulation requires jointly accounting for shared-replication uncertainty and deciding both where to sample and how many replications to allocate.

To address this joint challenge, we exploit the fact that the two objectives are estimated from the same raw simulator runs rather than treating them as separate noisy responses. A set of r replications at a design produces a single bivariate observation—the sample mean and sample variance—whose finite-sample uncertainty and cross-covariance both depend on r. The replication count is therefore not merely a tuning parameter: it simultaneously controls the joint estimation precision of the two objectives and the simulation cost. We derive this finite-replication covariance exactly and embed it in a hierarchical multi-output Gaussian process (GP) that separates dependence induced by shared estimation errors from dependence between the underlying mean and variance surfaces. Building on this posterior, hierarchical expected quantile hypervolume improvement (HEQHVI) treats the design and replication count as a joint decision and selects the pair with the largest expected improvement in the posterior-quantile Pareto front per simulator run. This allows the algorithm to balance exploring new designs against allocating additional replications to refine promising ones. For parallel evaluation, we extend HEQHVI with a minimum-energy design (MED) criterion that preserves the best serial candidate while spreading the remaining batch members across promising but distinct regions. Figure 1 summarizes the resulting workflow from shared-replication uncertainty modeling to cost-aware serial and parallel acquisition.

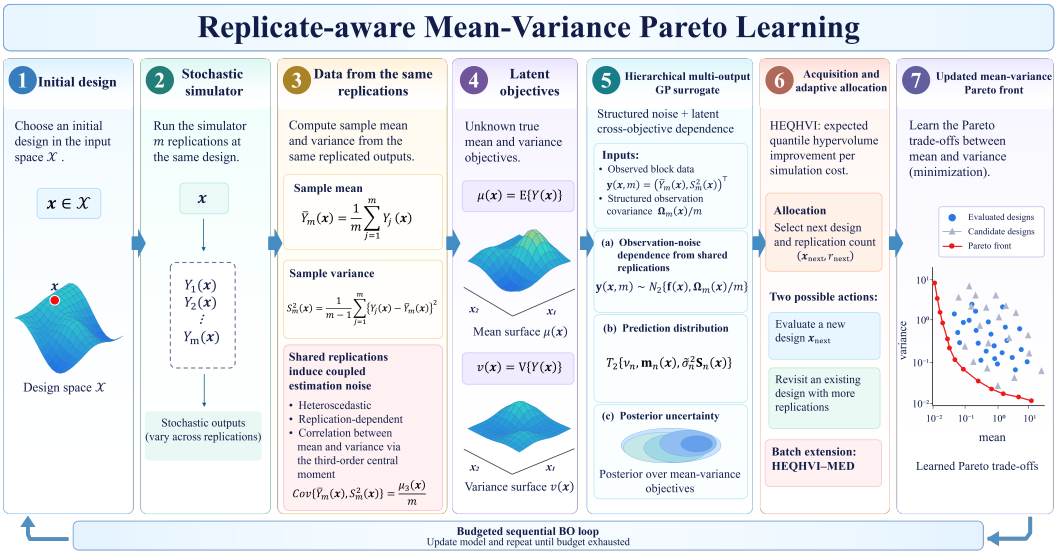


Figure 1: Replicate-aware Bayesian optimization framework for mean–variance Pareto learning.

2 REPLICATE-AWARE SURROGATE MODELING

We now construct the probabilistic surrogate that underlies the replicate-aware Pareto acquisition. The key is to treat the sample mean and sample variance computed from the same finite set of simulator runs as a joint noisy observation: we first characterize its replication-dependent sampling covariance and then embed this structure in a hierarchical multi-output GP for the latent mean and variance surfaces.

2.1 MEAN–VARIANCE OBJECTIVES AND SHARED-REPLICATION OBSERVATION MODEL

Let $\mathcal{X}$ be a compact design space, and let $Y(\boldsymbol{x})$ denote the scalar random output of one simulator run at $\boldsymbol{x} \in \mathcal{X}$. We characterize each design by its mean and variance,

$$\boldsymbol{f}(\boldsymbol{x}) = \begin{pmatrix} \mu(\boldsymbol{x}) \\ v(\boldsymbol{x}) \end{pmatrix} = \begin{pmatrix} \mathbb{E}\{Y(\boldsymbol{x})\} \\ \mathrm{Var}\{Y(\boldsymbol{x})\} \end{pmatrix}.$$

Under the minimization convention, $\boldsymbol{x}$ is Pareto optimal if there is no $\boldsymbol{x}' \in \mathcal{X}$ such that $\boldsymbol{f}(\boldsymbol{x}') \leq \boldsymbol{f}(\boldsymbol{x})$ componentwise, with at least one strict inequality. Let $\mathcal{X}^\star$ denote the set of Pareto-optimal designs and let $\mathcal{P}^\star = \{\boldsymbol{f}(\boldsymbol{x}) : \boldsymbol{x} \in \mathcal{X}^\star\}$ denote the corresponding Pareto front.

The latent objectives $\mu(\boldsymbol{x})$ and $v(\boldsymbol{x})$ are not observed directly. Instead, a finite set of simulator replications provides joint estimates of both quantities. For a fixed $\boldsymbol{x} \in \mathcal{X}$, let $Y_1(\boldsymbol{x}), \ldots, Y_m(\boldsymbol{x})$ be $m \geq 2$ independent copies of $Y(\boldsymbol{x})$, and define

$$\bar{Y}_m(\boldsymbol{x}) = \frac{1}{m}\sum_{j=1}^{m} Y_j(\boldsymbol{x}), \qquad S_m^2(\boldsymbol{x}) = \frac{1}{m-1}\sum_{j=1}^{m}\{Y_j(\boldsymbol{x}) - \bar{Y}_m(\boldsymbol{x})\}^2.$$

Let $\boldsymbol{y}(\boldsymbol{x}, m) = (\bar{Y}_m(\boldsymbol{x}), S_m^2(\boldsymbol{x}))^\top$, which is unbiased for $\boldsymbol{f}(\boldsymbol{x})$. Because both components are computed from the same raw simulator outputs, their estimation errors are generally dependent and change with the replication count. For $k \in \{3, 4\}$, define the central moment $\mu_k(\boldsymbol{x}) = \mathbb{E}[\{Y(\boldsymbol{x}) - \mu(\boldsymbol{x})\}^k]$. The following result makes this shared-replication structure explicit.

Proposition 1. *For any fixed $\boldsymbol{x} \in \mathcal{X}$ such that $\mathbb{E}|Y(\boldsymbol{x})|^4 < \infty$, and any $m \geq 2$,*

$$\begin{aligned} &\mathrm{Var}\{\bar{Y}_m(\boldsymbol{x})\} = \frac{v(\boldsymbol{x})}{m}, \quad \mathrm{Cov}\{\bar{Y}_m(\boldsymbol{x}), S_m^2(\boldsymbol{x})\} = \frac{\mu_3(\boldsymbol{x})}{m}, \\ &\mathrm{Var}\{S_m^2(\boldsymbol{x})\} = \frac{1}{m}\left\{\mu_4(\boldsymbol{x}) - \frac{m-3}{m-1} v(\boldsymbol{x})^2\right\}. \end{aligned} \tag{1}$$

Proposition 1 reveals two features that would be missed by treating the sample mean and sample variance as independent noisy objectives. First, their estimation errors are generally correlated, with cross-covariance determined by the third central moment $\mu_3(\boldsymbol{x})$. Second, the uncertainty of the sample variance depends on both the fourth central moment and the replication count. Thus, the joint observation noise depends on the local output distribution and changes as simulation effort is allocated. For symmetric output distributions the cross-covariance vanishes, while Gaussian sampling yields the stronger result that $\bar{Y}_m(\boldsymbol{x})$ and $S_m^2(\boldsymbol{x})$ are independent.

Although Proposition 1 gives the exact finite-sample covariance, the joint distribution of $(\bar{Y}_m(\boldsymbol{x}), S_m^2(\boldsymbol{x}))$ is generally not available in a convenient form. We therefore use the joint asymptotic distribution to obtain a tractable working model while retaining the exact finite-replication covariance.

Proposition 2. *Under the moment and regularity conditions stated in Appendix A, for each fixed $\boldsymbol{x} \in \mathcal{X}$,*

$$\sqrt{m}\{\boldsymbol{y}(\boldsymbol{x}, m) - \boldsymbol{f}(\boldsymbol{x})\} \Rightarrow \mathcal{N}_2\{0, \boldsymbol{\Omega}_\infty(\boldsymbol{x})\}, \tag{2}$$

where

$$\boldsymbol{\Omega}_\infty(\boldsymbol{x}) = \begin{pmatrix} v(\boldsymbol{x}) & \mu_3(\boldsymbol{x}) \\ \mu_3(\boldsymbol{x}) & \mu_4(\boldsymbol{x}) - v(\boldsymbol{x})^2 \end{pmatrix}.$$

Proposition 2 provides the asymptotic Gaussian shape, whereas Proposition 1 determines the covariance that should be used at a finite replication count. In particular, the asymptotic covariance misses a finite-m correction in the sample-variance component. Define

$$\boldsymbol{\Omega}_m(\boldsymbol{x}) = \boldsymbol{\Omega}_\infty(\boldsymbol{x}) + \mathrm{diag}\{0, 2v(\boldsymbol{x})^2/(m-1)\},$$

so that $\mathrm{Cov}\{\boldsymbol{y}(\boldsymbol{x}, m)\} = \boldsymbol{\Omega}_m(\boldsymbol{x})/m$ exactly. This leads to the moment-matched Gaussian working model

$$\boldsymbol{y}(\boldsymbol{x}, m) = \boldsymbol{f}(\boldsymbol{x}) + \varepsilon(\boldsymbol{x}, m), \qquad \varepsilon(\boldsymbol{x}, m) \dot{\sim} \mathcal{N}_2\{0, \boldsymbol{\Omega}_m(\boldsymbol{x})/m\}. \tag{3}$$

Thus, Equation (3) combines the exact finite-replication mean and covariance from Proposition 1 with the asymptotic Gaussian justification in Equation (2). In implementation, $\boldsymbol{\Omega}_m(\boldsymbol{x})$ is replaced by a plug-in estimate constructed from the available simulator outputs.

Let $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n$ be the distinct evaluated design points, with m_i cumulative replications available at $\boldsymbol{x}_i$. For a candidate pair $(\boldsymbol{x}, r)$ that collects r new replications at $\boldsymbol{x}$, let $\widehat{\boldsymbol{\Omega}}_{n,r}(\boldsymbol{x})$ denote the estimated scaled sampling covariance for the corresponding mean–variance observation. At an evaluated design, the available raw outputs are used to estimate the local quantities entering $\boldsymbol{\Omega}_m(\boldsymbol{x})$, and the candidate replication count r determines the finite-replication correction. At an unevaluated input, these covariance quantities are interpolated from the evaluated designs using matrix inverse-distance weighting (Osorio & Bierlaire, 2013). The full construction is given in Appendix A.3.

2.2 Hierarchical Multi-Output GP with Structured Observation Noise

The observation model in Equation (3) separates the latent mean–variance objectives from the finite-replication error induced by shared simulator runs. We now embed this structured observation error in a multi-output GP for the latent objectives. A key distinction is that dependence arises at two levels: the latent mean and variance surfaces may be correlated across the design space, while their estimation errors are correlated locally because they are computed from the same replications. Our surrogate models these two sources separately. Because the two objectives may have different numerical scales, we apply a fixed standardization determined from the initial design; a subscript s denotes quantities on this standardized scale, and predictions are transformed back to the original scale for Pareto acquisition and interpretation.

For the standardized latent surface, we place the multi-output GP model

$$\boldsymbol{f}_s(\boldsymbol{x}) = \mathbf{H}(\boldsymbol{x})\boldsymbol{\beta} + \sigma \boldsymbol{z}(\boldsymbol{x}), \qquad \boldsymbol{z} \sim \mathrm{MGP}\{\mathbf{0}, \mathbf{K}_\vartheta(\cdot, \cdot)\},$$

where $\mathbf{H}(\boldsymbol{x}) \in \mathbb{R}^{2 \times p_\beta}$ is a specified trend matrix, $\boldsymbol{\beta} \in \mathbb{R}^{p_\beta}$ contains its coefficients, $\sigma^2 > 0$ is the process-scale parameter, and ϑ contains the kernel parameters. The matrix-valued kernel $\mathbf{K}_\vartheta(\boldsymbol{x}, \boldsymbol{x}') \in \mathbb{R}^{2 \times 2}$ captures dependence between the latent objective surfaces (Svenson & Santner, 2016), distinct from the dependence in their finite-replication estimation errors. We use an intrinsic coregionalization model (ICM) (Alvarez et al., 2012).

Let $\mathbf{Q}_r(\boldsymbol{x})$ denote the normalized version, on the standardized scale, of the finite-r scaled sampling covariance matrix $\boldsymbol{\Omega}_r(\boldsymbol{x})$. We model the observation-error covariance from r replications as

$$\mathrm{Cov}\{\boldsymbol{y}_s(\boldsymbol{x}, r) \mid \boldsymbol{f}_s(\boldsymbol{x})\} = \frac{\sigma^2 \eta}{r} \mathbf{Q}_r(\boldsymbol{x}),$$

where $\eta > 0$ controls its magnitude relative to the latent-process variance. Let $\widehat{\mathbf{Q}}_{n,r}(\boldsymbol{x})$ denote the corresponding plug-in estimate obtained from the current data; its normalization and interpolation are detailed in Appendix A.3.

Define $\boldsymbol{y}_s = (\boldsymbol{y}_{s,1}^\top, \ldots, \boldsymbol{y}_{s,n}^\top)^\top \in \mathbb{R}^N$, where $N = 2n$, and let $\mathbf{H} \in \mathbb{R}^{N \times p_\beta}$ be the design matrix formed from $\mathbf{H}(\boldsymbol{x}_1), \ldots, \mathbf{H}(\boldsymbol{x}_n)$. Let $\mathbf{K}_n(\vartheta) = [\mathbf{K}_\vartheta(\boldsymbol{x}_i, \boldsymbol{x}_j)]_{i,j=1}^n \in \mathbb{R}^{N \times N}$ denote the latent GP covariance matrix. Assuming independent simulator randomness across distinct evaluated design points, the sampling-error covariance is block diagonal, giving

$$\mathbf{C}_n = \mathbf{K}_n(\vartheta) + \eta \,\mathrm{blockdiag}\left\{ \frac{\widehat{\mathbf{Q}}_{n,m_1}(\boldsymbol{x}_1)}{m_1}, \ldots, \frac{\widehat{\mathbf{Q}}_{n,m_n}(\boldsymbol{x}_n)}{m_n} \right\}.$$

This decomposition is central to the surrogate: $\mathbf{K}_n(\vartheta)$ captures dependence between the latent mean and variance surfaces across designs, whereas the block-diagonal term captures local, replication-dependent estimation uncertainty induced by the shared simulator runs. Hence,

$$\boldsymbol{y}_s \mid \boldsymbol{\beta}, \sigma^2 \sim \mathcal{N}_N\left(\mathbf{H}\boldsymbol{\beta}, \sigma^2 \mathbf{C}_n\right). \tag{4}$$

We estimate (ϑ, η) by profile likelihood and use generalized least squares for the trend; the fitting formulas are given in Appendix B.1 and updated after each new evaluation.

To propagate uncertainty in the trend and process scale, following Chen et al. (2024), we use $\pi(\boldsymbol{\beta}) \propto 1$ and $\sigma^2 \sim \mathrm{IG}(a, b)$, with $a, b > 0$. Let $\widehat{\boldsymbol{\beta}}_n$ be the generalized least-squares estimate of $\boldsymbol{\beta}$, and define

$$a_n = a + (N - p_\beta)/2, \qquad b_n = b + (\boldsymbol{y}_s - \mathbf{H}\widehat{\boldsymbol{\beta}}_n)^\top \mathbf{C}_n^{-1}(\boldsymbol{y}_s - \mathbf{H}\widehat{\boldsymbol{\beta}}_n)/2, \qquad \widetilde{\sigma}_n^2 = \frac{b_n}{a_n}.$$

Let $\mathcal{D}_n$ denote the observed sets of replications. For a test design $\boldsymbol{x}$, let $\boldsymbol{m}_n(\boldsymbol{x})$ and $\mathbf{S}_n(\boldsymbol{x})$ denote the generalized-kriging predictive location and scale component after integrating $\boldsymbol{\beta}$ and applying the fixed inverse standardization. Conditional on σ^2 and the fitted covariance quantities,

$$\boldsymbol{f}(\boldsymbol{x}) \mid \sigma^2, \mathcal{D}_n \sim \mathcal{N}_2\{\boldsymbol{m}_n(\boldsymbol{x}), \sigma^2 \mathbf{S}_n(\boldsymbol{x})\},$$

with detailed formulas given in Appendix B.2.

Proposition 3. *Suppose $N > p_\beta$, $\mathbf{H}$ has full column rank, $\mathbf{C}_n$ is positive definite, and $a, b > 0$. Conditional on the fitted covariance quantities in Equation (4), integrating $\boldsymbol{\beta}$ and σ^2 gives*

$$\sigma^2 \mid \mathcal{D}_n \sim \mathrm{IG}(a_n, b_n), \qquad \boldsymbol{f}(\boldsymbol{x}) \mid \mathcal{D}_n \sim T_2\big\{\nu_n, \boldsymbol{m}_n(\boldsymbol{x}), \widetilde{\sigma}_n^2 \mathbf{S}_n(\boldsymbol{x})\big\}, \quad \nu_n = 2a_n, \tag{5}$$

where $T_2(\nu, \mu, s)$ denotes a bivariate non-standardized Student-t distribution with ν degrees of freedom, location parameter μ, and scale parameter s.

Equation (5) propagates uncertainty in the trend and process scale in addition to the heterogeneous finite-replication observation error, and provides the posterior used by the Pareto acquisition criterion developed next. We use an unconstrained GP for the variance surface to preserve the additive joint observation model in Equation (3); consequently, occasional negative predictive values should be interpreted as surrogate artifacts rather than physical variances.

3 Replicate-Aware Pareto Acquisition

The surrogate in Section 2.2 provides a posterior distribution for the latent mean–variance objectives while accounting for shared finite-replication uncertainty. We use this posterior to choose a design–replication pair $(\boldsymbol{x}, r)$, where $\boldsymbol{x}$ determines where to evaluate and r controls both the precision of the posterior update and the simulation cost. Previously evaluated designs remain eligible, so the acquisition can trade off exploring new designs against refining existing ones. We first develop the serial criterion and then extend it to parallel evaluation. Throughout, n denotes the number of distinct evaluated designs, whereas t indexes the BO iteration; these can differ because designs may be reevaluated.

3.1 Serial HEQHVI

To score a candidate pair $(\boldsymbol{x}, r)$ before observing its simulator outputs, we measure its prospective improvement to the posterior-quantile Pareto front. Following expected quantile improvement (Picheny et al., 2013), fix a marginal quantile level $\tau \in (0.5, 1)$ and let $c_{\tau,n} = T_{\nu_n}^{-1}(\tau)$, where T_{ν_n} is the CDF of a standard Student-t distribution with ν_n degrees of freedom. Because both objectives are minimized, an upper marginal quantile penalizes posterior uncertainty. Using Proposition 3, define

$$\boldsymbol{\xi}_n(\boldsymbol{x}) = \boldsymbol{m}_n(\boldsymbol{x}) + c_{\tau,n}\widetilde{\sigma}_n \begin{pmatrix} \sqrt{[\mathbf{S}_n(\boldsymbol{x})]_{11}} \\ \sqrt{[\mathbf{S}_n(\boldsymbol{x})]_{22}} \end{pmatrix}, \tag{6}$$

where $[\mathbf{S}_n(\boldsymbol{x})]_{jj}$ is the (j, j) entry of $\mathbf{S}_n(\boldsymbol{x})$, $j = 1, 2$.

Let $\mathcal{R} \subset \{2, 3, \ldots\}$ be the finite set of permitted replication counts. A future batch of r replications at $\boldsymbol{x}$ produces $\widetilde{\boldsymbol{y}}(\boldsymbol{x}, r) = (\bar{Y}_r^{\mathrm{new}}(\boldsymbol{x}), (S_r^{\mathrm{new}}(\boldsymbol{x}))^2)^\top$, with $\widetilde{\boldsymbol{y}}(\boldsymbol{x}, r) \mid \boldsymbol{f}(\boldsymbol{x}), \sigma^2, \mathcal{D}_n \sim \mathcal{N}_2\{\boldsymbol{f}(\boldsymbol{x}), \sigma^2 \boldsymbol{\Lambda}_r(\boldsymbol{x})\}$, where $\boldsymbol{\Lambda}_r(\boldsymbol{x})$ is detailed in Appendix C.1. Writing $\mathbf{S} = \mathbf{S}_n(\boldsymbol{x})$ and $\boldsymbol{\Lambda} = \boldsymbol{\Lambda}_r(\boldsymbol{x})$, Gaussian conditioning gives

$$\begin{aligned} \boldsymbol{m}_+(\boldsymbol{x}, r) &= \boldsymbol{m}_n(\boldsymbol{x}) + \mathbf{S}(\mathbf{S} + \boldsymbol{\Lambda})^{-1}\{\widetilde{\boldsymbol{y}}(\boldsymbol{x}, r) - \boldsymbol{m}_n(\boldsymbol{x})\}, \\ \mathbf{S}_+(\boldsymbol{x}, r) &= \mathbf{S} - \mathbf{S}(\mathbf{S} + \boldsymbol{\Lambda})^{-1}\mathbf{S}. \end{aligned}$$

Thus, r enters through $\boldsymbol{\Lambda}_r(\boldsymbol{x})$ and controls both the possible posterior-location update and the uncertainty remaining after evaluation.

Because the future observation is unknown, the updated quantile vector is itself random. Its distribution under the current fitted surrogate is as follows.

Proposition 4. *Holding the current ν_n and $\widetilde{\sigma}_n$ fixed, the updated marginal-quantile vector is*

$$\boldsymbol{\Xi}_+(\boldsymbol{x}, r) = \boldsymbol{m}_+(\boldsymbol{x}, r) + c_{\tau,n}\widetilde{\sigma}_n \begin{pmatrix} \sqrt{[\mathbf{S}_+(\boldsymbol{x}, r)]_{11}} \\ \sqrt{[\mathbf{S}_+(\boldsymbol{x}, r)]_{22}} \end{pmatrix}.$$

Conditional on $\mathcal{D}_n$ and the fitted covariance quantities,

$$\boldsymbol{\Xi}_+(\boldsymbol{x}, r) \mid \mathcal{D}_n \sim T_2\left\{\nu_n,\, \boldsymbol{m}_n(\boldsymbol{x}) + c_{\tau,n}\widetilde{\sigma}_n \begin{pmatrix} \sqrt{[\mathbf{S}_+(\boldsymbol{x}, r)]_{11}} \\ \sqrt{[\mathbf{S}_+(\boldsymbol{x}, r)]_{22}} \end{pmatrix},\, \widetilde{\sigma}_n^2\{\mathbf{S}_n(\boldsymbol{x}) - \mathbf{S}_+(\boldsymbol{x}, r)\}\right\}. \quad (7)$$

Proposition 4 turns each $(\boldsymbol{x}, r)$ into a distribution over post-evaluation quantile vectors, allowing its Pareto gain to be averaged over the unseen simulator outcome. The fitted covariance parameters are held fixed in this prospective calculation and re-estimated after the new outputs are obtained.

Let the nondominated vectors among $\{\boldsymbol{\xi}_n(\boldsymbol{x}_i) : i = 1, \ldots, n\}$ define the current quantile-based Pareto front $\mathcal{P}_n^{(\tau)}$. For a prospective quantile vector $\boldsymbol{z}$ at $\boldsymbol{x}$, let $\mathcal{P}_+(\boldsymbol{x}, \boldsymbol{z})$ be the front obtained by adding $\boldsymbol{z}$ if $\boldsymbol{x}$ is new, or replacing its current quantile vector if $\boldsymbol{x}$ is reevaluated, followed by nondominance filtering. Let $\boldsymbol{r}_{\mathrm{ref}}$ be a fixed reference point with both coordinates worse than the objective values of interest. Assuming $\nu_n > 2$, define

$$\mathrm{HEQHVI}_n(\boldsymbol{x}, r) = \mathbb{E}\left[\left\{\mathrm{HV}(\mathcal{P}_+\{\boldsymbol{x}, \boldsymbol{\Xi}_+(\boldsymbol{x}, r)\}; \boldsymbol{r}_{\mathrm{ref}}) - \mathrm{HV}\left(\mathcal{P}_n^{(\tau)}; \boldsymbol{r}_{\mathrm{ref}}\right)\right\}_+ \middle| \mathcal{D}_n\right]. \quad (8)$$

Here, HV denotes hypervolume. HEQHVI measures the expected positive hypervolume change induced by the prospective posterior update, rather than scoring the noisy mean–variance observation directly. For reevaluation, the current vector is replaced rather than added, and the positive part credits only net improvement. The equivalent HVI representation and exact pooling update are given in Appendices C.3 and C.2.

A batch of r replications costs r simulator runs, so we select

$$(\boldsymbol{x}_{\mathrm{next}}, r_{\mathrm{next}}) \in \arg\max_{\boldsymbol{x}\in\mathcal{X},\, r\in\mathcal{R}} \frac{\mathrm{HEQHVI}_n(\boldsymbol{x}, r)}{r}. \quad (9)$$

Equation (9) makes replication allocation endogenous: additional replications are selected only when their expected Pareto benefit justifies their additional cost. Because evaluated designs remain feasible, the same criterion chooses between exploration and refinement. With fixed plug-in moments, increasing r reduces the residual conditional covariance, although the cost-adjusted acquisition need not be monotone in r; see Appendix C.4. Numerical optimization is detailed in Appendix D.1.

3.2 Parallel HEQHVI–MED

Equation (9) selects a single design–replication pair. When q evaluations can be run in parallel, however, selecting several high-valued serial candidates can concentrate the batch in one region and waste parallel capacity through redundant information. We therefore preserve the best serial action and use MED to diversify the remaining design locations (Li et al., 2023). MED treats design points as charged particles and favors configurations with low maximum pairwise energy (Joseph et al., 2015):

$$E(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_q) = \max_{i<j} \frac{\mathcal{Q}_n(\boldsymbol{x}_i)\mathcal{Q}_n(\boldsymbol{x}_j)}{\|\boldsymbol{x}_i - \boldsymbol{x}_j\|_2}, \quad (10)$$

where $\mathcal{Q}_n(\boldsymbol{x}) > 0$ is the charge at $\boldsymbol{x}$. All MED distances are computed in the scaled input coordinates used by the GP.

Because HEQHVI scores design–replication pairs whereas MED diversifies design locations, we first profile out the replication decision. Define $\alpha_n(\boldsymbol{x}, r) = \mathrm{HEQHVI}_n(\boldsymbol{x}, r)/r$ and

$$\bar{\alpha}_n(\boldsymbol{x}) = \max_{r\in\mathcal{R}} \alpha_n(\boldsymbol{x}, r). \quad (11)$$

Thus, $\bar{\alpha}_n(\boldsymbol{x})$ is the best cost-adjusted serial value attainable at $\boldsymbol{x}$, and once $\boldsymbol{x}$ is selected its replication count is chosen from the maximizers in Equation (11). For $\bar{\alpha}_n(\boldsymbol{x}) > 0$, set $\mathcal{Q}_n(\boldsymbol{x}) = 1/\bar{\alpha}_n(\boldsymbol{x})$. Substituting these charges into Equation (10) gives the equivalent criterion

$$\max_{\boldsymbol{x}_1,\ldots,\boldsymbol{x}_q\in\mathcal{X}} \min_{i<j} \bar{\alpha}_n(\boldsymbol{x}_i)\bar{\alpha}_n(\boldsymbol{x}_j)\|\boldsymbol{x}_i - \boldsymbol{x}_j\|_2. \quad (12)$$

Hence, a batch is favored when its members are both individually valuable under HEQHVI and well separated in the design space.

Direct joint optimization of Equation (12) can be difficult for a large batch, so we construct the batch sequentially (Joseph et al., 2019). We anchor the batch at the serial optimum, $\tilde{\boldsymbol{x}}_1 \in \arg\max_{\boldsymbol{x}\in\mathcal{X}} \bar{\alpha}_n(\boldsymbol{x})$, and diversify the remaining members against both past evaluations and points already selected for the current batch. Let $\mathcal{H}_n = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n\}$ denote the previously evaluated designs. For a candidate $\boldsymbol{x}$, define its distance to the historical designs as

$$d_{\mathcal{H},n}(\boldsymbol{x}) = \min_{\boldsymbol{z}\in\mathcal{H}_n} \|\boldsymbol{x} - \boldsymbol{z}\|_2.$$

For the previously selected members $\mathcal{B}_{j-1} = \{\tilde{\boldsymbol{x}}_1, \ldots, \tilde{\boldsymbol{x}}_{j-1}\}$ of the current batch, define the acquisition-weighted separation

$$d_{\mathcal{B},n,j}(\boldsymbol{x}) = \min_{\ell<j} \bar{\alpha}_n(\tilde{\boldsymbol{x}}_\ell)\|\boldsymbol{x} - \tilde{\boldsymbol{x}}_\ell\|_2.$$

The historical distance discourages repeated coverage of previously evaluated regions. The acquisition weight in $d_{\mathcal{B},n,j}$ is included so that multiplication by $\bar{\alpha}_n(\boldsymbol{x})$ recovers the pairwise factor in Equation (12). For $j = 2, \ldots, q$, choose

$$\tilde{\boldsymbol{x}}_j \in \arg\max_{\boldsymbol{x}\in\mathcal{X}} \min\left\{\bar{\alpha}_n(\boldsymbol{x}) d_{\mathcal{H},n}(\boldsymbol{x}),\ \bar{\alpha}_n(\boldsymbol{x}) d_{\mathcal{B},n,j}(\boldsymbol{x})\right\}. \tag{13}$$

The acquisition factor favors promising locations, while the two separation terms discourage redundancy with historical evaluations and within the current batch. Each selected $\tilde{\boldsymbol{x}}_j$ is then paired with a replication count attaining the maximum in Equation (11). HEQHVI–MED therefore preserves the best serial decision while using the remaining parallel evaluations to cover distinct high-value regions.

4 Experiments

We evaluate serial HEQHVI and parallel HEQHVI–MED on four synthetic mean–variance scenarios and an F119 aero-engine application. The synthetic experiments assess Pareto-front recovery against known objectives; the engine study uses independent simulations to evaluate selected operating conditions. Additional experimental details are given in Appendix F.

4.1 Synthetic Benchmarks

We use the two-dimensional Branin and six-dimensional Hartmann benchmarks with stochastic output $Y(\boldsymbol{x}) = G(\boldsymbol{x}) + s_Y(\boldsymbol{x})\varepsilon$, where G is the deterministic response, s_Y is the input-dependent output standard deviation, and ε has zero mean and unit variance. We set s_Y to be largest near the minimizers of G. Designs with favorable mean responses therefore have relatively large output variances, producing a nontrivial mean–variance Pareto front. Gaussian and standardized Gamma errors yield the same latent mean–variance objectives but different higher-order sampling behavior, including zero versus nonzero covariance between the sample mean and sample variance. Combining the two benchmarks and two error distributions yields four scenarios.

We compare adaptive and fixed-replication serial HEQHVI and parallel HEQHVI–MED with GPS-C (Wang et al., 2025), the sum of mean and variance (SMV) applied in Makarova et al. (2021), and serial NEHVI and parallel qNEHVI (Daulton et al., 2020; 2021). Each scenario is evaluated over 100 independent optimization runs, with 50 BO iterations following a shared initial design. We report performance against BO iterations to assess Pareto-front recovery with a limited number of sequential decision stages. Adaptive HEQHVI selects $r \in \{5, 10, 15\}$, whereas fixed-replication HEQHVI and HEQHVI–MED use $r = 5$. Both parallel policies use five design points per iteration and five replications per design point. At every iteration, we form the nondominated set from the true mean and variance at the distinct design points evaluated by each method. Performance is measured by dominated hypervolume and inverted generational distance (IGD$^+$) relative to a continuous-domain reference front constructed offline.

Figure 2 presents the two-dimensional Branin results, with Gaussian errors in the first row and standardized Gamma errors in the second. NEHVI/qNEHVI and our methods are the two leading approaches under the matched serial and parallel settings. Among the serial HEQHVI policies,

adaptive replication yields higher hypervolume and lower IGD$^+$ than fixed replication under both output laws, suggesting a benefit from selecting r during optimization; the difference is more pronounced under Gaussian errors. Both serial HEQHVI policies also finish with higher hypervolume and lower IGD$^+$ than serial NEHVI. With five design points evaluated per iteration, HEQHVI–MED attains higher hypervolume and lower IGD$^+$ than the serial HEQHVI policies within the same number of BO iterations, indicating that the batch construction effectively uses the available parallel evaluations. Compared with qNEHVI, HEQHVI–MED improves hypervolume and reduces IGD$^+$ more rapidly during the early iterations. It finishes with slightly higher hypervolume under both output laws and lower IGD$^+$ under Gamma errors; their final IGD$^+$ values are nearly identical under Gaussian errors. The corresponding Hartmann results are reported in Appendix F.

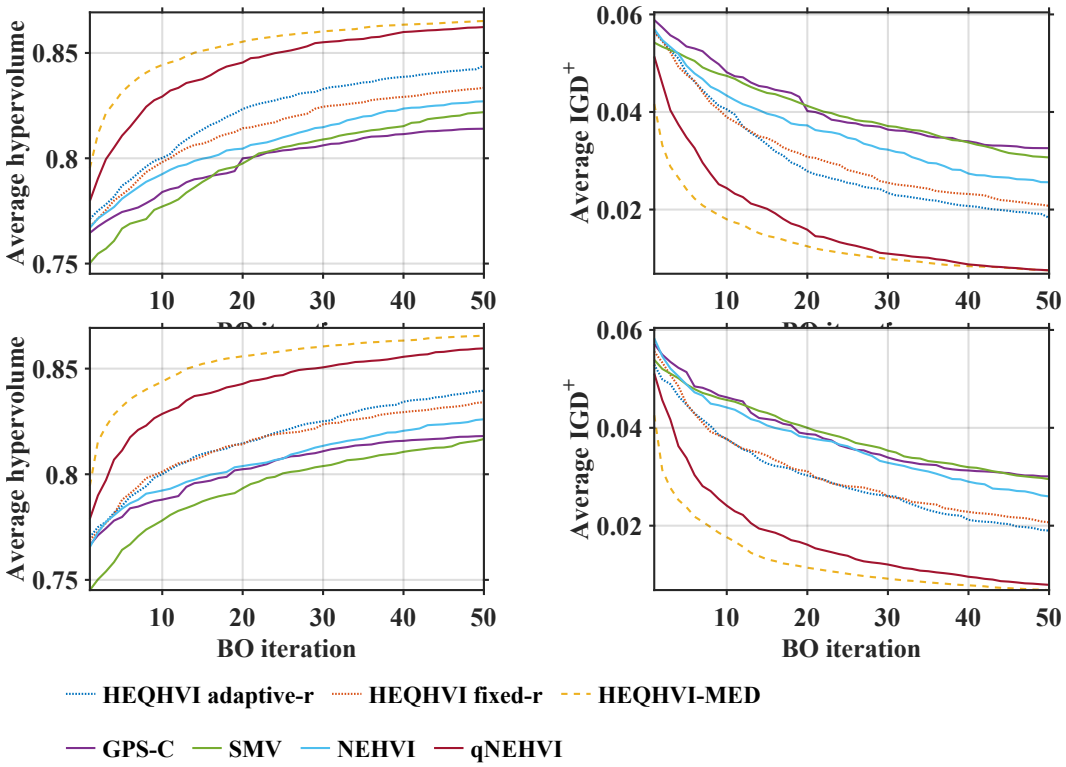


Figure 2: Performance on the two-dimensional Branin benchmark averaged over 100 independent optimization runs. The first row uses Gaussian errors and the second uses Gamma errors. The left column shows the average hypervolume, where higher is better, and the right column shows IGD$^+$, where lower is better. Shaded regions denote 95% confidence intervals.

4.2 F119 Aero-Engine Optimization

We consider operating-condition selection for an F119 turbofan using a dynamic MATLAB/Simulink model built with T-MATS (Chapman et al., 2014a). The model couples thermodynamic components, turbomachinery performance maps, shaft dynamics, and feedback control through iterative numerical solvers. This nonlinear simulation is treated as an expensive black box. Figure 4 in Appendix F.2 shows the simulation architecture. According to Chapman et al. (2014b), T-MATS supports stochastic simulation through sensor and fuel-actuator models. The decision variables are flight altitude and Mach number, $\boldsymbol{x} = (H, M)^\top \in [6, 12] \times [0.5, 1.0]$. With $Y(\boldsymbol{x}) = -T(\boldsymbol{x})$, the objectives are $\mathbb{E}\{Y(\boldsymbol{x})\} = -\mathbb{E}\{T(\boldsymbol{x})\}$ and $\mathrm{Var}\{Y(\boldsymbol{x})\} = \mathrm{Var}\{T(\boldsymbol{x})\}$. Thus, minimizing $(\mathbb{E}\{Y(\boldsymbol{x})\}, \mathrm{Var}\{Y(\boldsymbol{x})\})$ is equivalent to maximizing mean thrust while minimizing thrust variability.

Figure 3 shows that adaptive HEQHVI provides the strongest serial performance among the proposed variants, ending with higher hypervolume and lower IGD$^+$ than fixed-r HEQHVI. Fixed-r HEQHVI nevertheless outperforms NEHVI, indicating that the improvement is not solely attributable to adaptive replication. The advantage is most pronounced in the parallel setting: HEQHVI–MED improves

more rapidly and maintains both higher hypervolume and lower IGD^{+} than qNEHVI. Overall, the trajectories show improved final Pareto-front recovery for the proposed serial and parallel variants under the corresponding experimental settings.

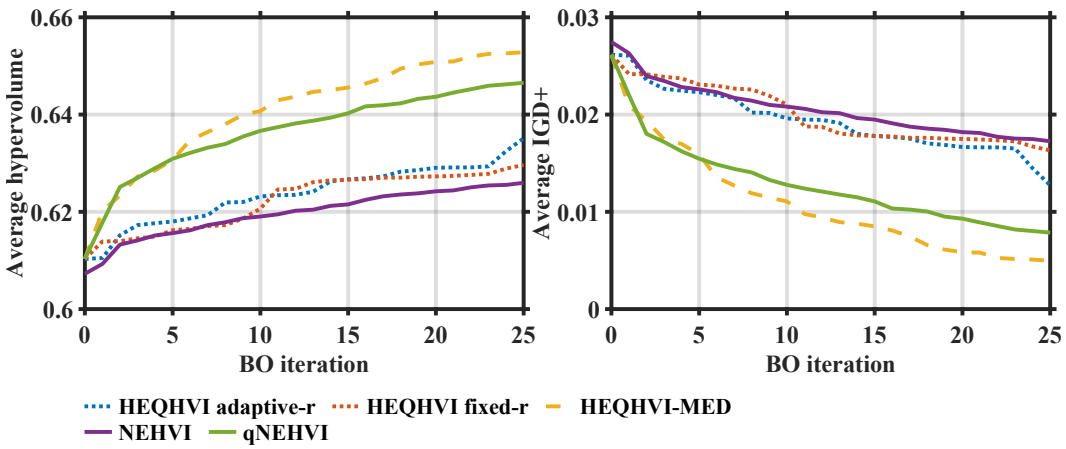


Figure 3: Average hypervolume (left) and IGD^{+} (right) over 25 BO iterations for the F119 aero-engine case. Larger hypervolume and smaller IGD^{+} indicate better mean–variance Pareto-front approximations.

5 Conclusion

We developed a replicate-aware Bayesian optimization framework for learning mean–variance Pareto trade-offs from stochastic simulators. The key observation is that the sample mean and sample variance are computed from the same finite set of simulator runs, so their estimation errors are coupled and depend on the replication count. We characterized this finite-replication covariance and incorporated it into a hierarchical multi-output GP that separates sampling-error dependence from dependence between the latent mean and variance surfaces. Building on this posterior, HEQHVI makes replication allocation part of the optimization decision by selecting a design–replication pair according to its expected improvement in the posterior-quantile Pareto front per simulator run. This allows the method to balance exploration of new designs against refinement of existing ones. HEQHVI–MED extends the same principle to parallel evaluation by preserving the best serial action while diversifying the remaining batch members. Synthetic experiments under symmetric and skewed output laws and an F119 aero-engine study illustrate the framework in both serial and parallel settings.

Several extensions remain open. One direction is to establish finite-budget guarantees for joint design and replication allocation with reevaluation. Another is to extend the framework to constrained stochastic optimization, where uncertainty in feasibility must be incorporated alongside the mean–variance objectives. For parallel simulation, jointly choosing the batch size and member-specific replication counts under a total budget could further balance spatial coverage, estimation precision, and the frequency of surrogate updates.

AI Use Statement

Generative AI tools were used to assist with literature organization, drafting and revising portions of the manuscript, language editing, LaTeX restructuring, and exploratory checks of mathematical derivations. AI-generated suggestions were treated as non-authoritative. The authors independently verified the mathematical claims and derivations, checked the cited literature against the original sources, and reviewed all AI-assisted text. Generative AI was not used to fabricate or alter experimental data or empirical results. The authors take full responsibility for the final content of this work, including all text, claims, code, and reported results.

REFERENCES

Mauricio A. Alvarez, Lorenzo Rosasco, and Neil D. Lawrence. Kernels for vector-valued functions: A review. *Foundations and Trends in Machine Learning*, 4(3):195–266, 2012.

Bruce E. Ankenman, Barry L. Nelson, and Jeremy Staum. Stochastic kriging for simulation metamodeling. *Operations Research*, 58(2):371–382, 2010.

Mickaël Binois and Jeffrey Larson. Adaptive replication strategies in trust-region-based bayesian optimization of stochastic functions. *Optimization Methods and Software*, pp. 1–34, 2026.

Mickaël Binois, Robert B. Gramacy, and Michael Ludkovski. Practical heteroskedastic gaussian process modeling for large simulation experiments. *Journal of Computational and Graphical Statistics*, 27(4):808–821, 2018.

Mickaël Binois, Jiangeng Huang, Robert B. Gramacy, and Michael Ludkovski. Replication or exploration? sequential design for stochastic simulation experiments. *Technometrics*, 61(1):7–23, 2019.

Sait Cakmak, Raul Astudillo Marban, Peter I. Frazier, and Enlu Zhou. Bayesian optimization of risk measures. In *Advances in Neural Information Processing Systems*, volume 33, pp. 20130–20141. Curran Associates, Inc., 2020.

Jeffryes W. Chapman, Thomas M. Lavelle, Jonathan S. Litt, and Ten-Huei Guo. A process for the creation of T-MATS propulsion system models from NPSS data. Technical Report NASA/TM-2014-218409, NASA Glenn Research Center, 2014a.

Jeffryes W. Chapman, Thomas M. Lavelle, Ryan D. May, Jonathan S. Litt, and Ten-Huei Guo. Propulsion system simulation using the toolbox for the modeling and analysis of thermodynamic systems (T-MATS). Technical Report NASA/TM-2014-218410, NASA Glenn Research Center, November 2014b.

Zhehui Chen, Simon Mak, and C. F. Jeff Wu. A hierarchical expected improvement method for bayesian optimization. *Journal of the American Statistical Association*, 119(546):1619–1632, 2024.

Zhongxiang Dai, Quoc Phong Nguyen, Sebastian Tay, Daisuke Urano, Richalynn Leong, Bryan Kian Hsiang Low, and Patrick Jaillet. Batch Bayesian optimization for replicable experimental design. *Advances in Neural Information Processing Systems*, 36:36476–36506, 2023.

Samuel Daulton, Maximilian Balandat, and Eytan Bakshy. Differentiable expected hypervolume improvement for parallel multi-objective bayesian optimization. In *Advances in Neural Information Processing Systems*, volume 33, pp. 9851–9864. Curran Associates, Inc., 2020.

Samuel Daulton, Maximilian Balandat, and Eytan Bakshy. Parallel Bayesian optimization of multiple noisy objectives with expected hypervolume improvement. *Advances in Neural Information Processing Systems*, 34:2187–2200, 2021.

Ryan-Rhys Griffiths, Alexander A. Aldrick, Miguel Garcia-Ortegon, Vidhi Lalchand, and Alpha A. Lee. Achieving robustness to aleatoric uncertainty with heteroscedastic bayesian optimisation. *Machine Learning: Science and Technology*, 3(1):015004, 2022.

Shogo Iwazaki, Yu Inatsu, and Ichiro Takeuchi. Mean-variance analysis in bayesian optimization under uncertainty. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pp. 973–981. PMLR, 2021.

Hamed Jalali, Inneke Van Nieuwenhuyse, and Victor Picheny. Comparison of kriging-based algorithms for simulation optimization with heterogeneous noise. *European Journal of Operational Research*, 261(1):279–301, 2017.

V. Roshan Joseph, Tirthankar Dasgupta, Rui Tuo, and C. F. Jeff Wu. Sequential exploration of complex surfaces using minimum energy designs. *Technometrics*, 57(1):64–74, 2015.

V. Roshan Joseph, Dianpeng Wang, Li Gu, Shiji Lyu, and Rui Tuo. Deterministic sampling of expensive posteriors using minimum energy designs. *Technometrics*, 61(3):297–308, 2019.

Joshua Knowles. Parego: A hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation*, 10(1): 50–66, 2006.

Miguel A. Lejeune and Wenbo Ma. Drone-delivery network for opioid overdose: Nonlinear integer queueing-optimization models and methods. *Operations Research*, 73(1):86–108, 2025.

ShiXiang Li, Yubin Tian, and Dianpeng Wang. Parallel efficient global optimization by using the minimum energy criterion. *Journal of Statistical Computation and Simulation*, 93(17):3104–3125, 2023.

Anastasia Makarova, Ilnura Usmanova, Ilija Bogunovic, and Andreas Krause. Risk-averse heteroscedastic Bayesian optimization. *Advances in Neural Information Processing Systems*, 34: 17235–17245, 2021.

Carolina Osorio and Michel Bierlaire. A simulation-based optimization framework for urban transportation problems. *Operations Research*, 61(6):1333–1345, 2013.

Victor Picheny, David Ginsbourger, Yann Richet, and Gregory Caplin. Quantile-based optimization of noisy computer experiments with tunable precision. *Technometrics*, 55(1):2–13, 2013.

Meng Qi, Yuanyuan Shi, Yongzhi Qi, Chenxin Ma, Rong Yuan, Di Wu, and Zuo-Jun Shen. A practical end-to-end inventory management model with deep learning. *Management Science*, 69 (2):759–773, 2023.

Daria Semochkina, Alexander I. J. Forrester, and David C. Woods. Multiobjective optimization using expected quantile improvement for decision making in disease outbreaks. *SIAM/ASA Journal on Uncertainty Quantification*, 13(1):228–250, 2025.

Joshua Svenson and Thomas J. Santner. Multiobjective optimization of expensive-to-evaluate deterministic computer simulator models. *Computational Statistics & Data Analysis*, 94:250–264, 2016.

Sebastian Shenghong Tay, Chuan Sheng Foo, Daisuke Urano, Richalynn Leong, and Bryan Kian Hsiang Low. A unified framework for Bayesian optimization under contextual uncertainty. In *International Conference on Learning Representations*, volume 2024, pp. 5913–5945, 2024.

Xiuxian Wang, L. Jeff Hong, Zhibin Jiang, and Haihui Shen. Gaussian process-based random search for continuous optimization via simulation. *Operations Research*, 73(1):385–407, 2025.

Jun Zhou, Mei Yang, Yong Zhan, and Li Xu. Digital twin application for reinforcement learning based optimal scheduling and reliability management enhancement of systems. *Solar Energy*, 252: 29–38, 2023.

A Finite-Replication Observation Model

Assumption 1. *For some $\delta > 0$,*

$$\sup_{\boldsymbol{x}\in\mathcal{X}} \mathbb{E}|Y(\boldsymbol{x})|^{4+\delta} < \infty.$$

Moreover, $v(\boldsymbol{x})$ is bounded away from zero, and $\mu(\boldsymbol{x})$, $v(\boldsymbol{x})$, $\mu_3(\boldsymbol{x})$, and $\mu_4(\boldsymbol{x})$ are continuous on the compact design space $\mathcal{X}$.

The uniform moment condition provides the finite second moments needed for the joint central limit theorem and the laws of large numbers used to estimate the spatially varying sampling covariance. The positive lower bound on $v(\boldsymbol{x})$ prevents degeneracy, while continuity of the moment functions supports spatial modeling over $\mathcal{X}$.

A.1 Proof of Proposition 1

Fix $\boldsymbol{x}$ and suppress it from the notation. Let $X_j = Y_j - \mu$, so that $\mathbb{E}(X_j) = 0$, $\mathbb{E}(X_j^2) = v$, $\mathbb{E}(X_j^3) = \mu_3$, and $\mathbb{E}(X_j^4) = \mu_4$. Also, define $\bar{X} = m^{-1}\sum_{j=1}^m X_j$. Then $\bar{Y}_m - \mu = \bar{X}$, and the unbiased sample variance can be written as

$$S_m^2 = \frac{1}{m-1}\left(\sum_{j=1}^m X_j^2 - m\bar{X}^2\right).$$

To establish Equation (1), it suffices to calculate the three distinct entries $\mathrm{Var}(\bar{Y}_m)$, $\mathrm{Cov}(\bar{Y}_m, S_m^2)$, and $\mathrm{Var}(S_m^2)$.

We first consider the variance of the sample mean. Since the X_j's are independent and identically distributed,

$$\mathrm{Var}(\bar{Y}_m) = \mathrm{Var}\left(\frac{1}{m}\sum_{j=1}^m X_j\right) = \frac{1}{m^2}\sum_{j=1}^m \mathrm{Var}(X_j) = \frac{v}{m}.$$

We next calculate the covariance between the sample mean and the sample variance. Since $\mathbb{E}(\bar{X}) = 0$ and $\mathbb{E}(S_m^2) = v$, it follows that

$$\mathrm{Cov}(\bar{Y}_m, S_m^2) = \mathbb{E}(\bar{X}S_m^2) = \frac{1}{m-1}\left[\mathbb{E}\left\{\bar{X}\sum_{j=1}^m X_j^2\right\} - \mathbb{E}(m\bar{X}^3)\right].$$

We evaluate the two expectations separately. For the first one,

$$\mathbb{E}\left\{\bar{X}\sum_{j=1}^m X_j^2\right\} = \frac{1}{m}\sum_{k=1}^m\sum_{j=1}^m \mathbb{E}(X_k X_j^2) = \frac{1}{m}\left[\sum_{j=1}^m \mathbb{E}(X_j^3) + \sum_{k\neq j}\mathbb{E}(X_k X_j^2)\right].$$

When $k \neq j$, independence and centering give $\mathbb{E}(X_k X_j^2) = \mathbb{E}(X_k)\mathbb{E}(X_j^2) = 0$. Therefore,

$$\mathbb{E}\left\{\bar{X}\sum_{j=1}^m X_j^2\right\} = \frac{1}{m}\sum_{j=1}^m \mu_3 = \mu_3.$$

For the second expectation, expanding the third power gives

$$\mathbb{E}\left(\sum_{j=1}^m X_j\right)^3 = \sum_{j=1}^m \mathbb{E}(X_j^3) = m\mu_3,$$

because every cross-product contains at least one centered variable appearing only once. Hence,

$$\mathbb{E}(m\bar{X}^3) = \frac{1}{m^2}\mathbb{E}\left(\sum_{j=1}^m X_j\right)^3 = \frac{\mu_3}{m}.$$

Combining these two calculations, we obtain

$$\mathrm{Cov}(\bar{Y}_m, S_m^2) = \frac{1}{m-1}\left(\mu_3 - \frac{\mu_3}{m}\right) = \frac{\mu_3}{m}.$$

It remains to calculate the variance of the sample variance. Define

$$Q = \sum_{j=1}^m X_j^2, \qquad T = \sum_{j=1}^m X_j.$$

Then $S_m^2 = (Q - T^2/m)/(m-1)$, and therefore

$$\mathbb{E}\{(S_m^2)^2\} = \frac{1}{(m-1)^2}\left[\mathbb{E}(Q^2) - \frac{2}{m}\mathbb{E}(QT^2) + \frac{1}{m^2}\mathbb{E}(T^4)\right].$$

We now calculate the three expectations on the right-hand side. First, expanding Q^2 yields

$$\begin{aligned}\mathbb{E}(Q^2) &= \mathbb{E}\left[\sum_{j=1}^{m} X_j^4 + 2\sum_{1\le j<k\le m} X_j^2 X_k^2\right] \\ &= \sum_{j=1}^{m}\mathbb{E}(X_j^4) + 2\sum_{1\le j<k\le m}\mathbb{E}(X_j^2)\mathbb{E}(X_k^2) \\ &= m\mu_4 + m(m-1)v^2.\end{aligned}$$

Second, because $T^2 = Q + 2\sum_{1\le j<k\le m} X_j X_k$, we have

$$QT^2 = Q^2 + 2\sum_{i=1}^{m}\sum_{1\le j<k\le m} X_i^2 X_j X_k.$$

Every term in the double sum has expectation zero. Indeed, if $i \notin \{j,k\}$, both X_j and X_k appear only once; if $i = j$ or $i = k$, the remaining variable still appears only once. Independence and centering therefore imply $\mathbb{E}(QT^2) = \mathbb{E}(Q^2) = m\mu_4 + m(m-1)v^2$.

Third, in the expansion of T^4, a term has a nonzero expectation only when its indices occur either four times or in two pairs. Consequently,

$$\begin{aligned}\mathbb{E}(T^4) &= \sum_{j=1}^{m}\mathbb{E}(X_j^4) + 6\sum_{1\le j<k\le m}\mathbb{E}(X_j^2 X_k^2) \\ &= m\mu_4 + 6\binom{m}{2}v^2 \\ &= m\mu_4 + 3m(m-1)v^2.\end{aligned}$$

Substituting these three results into the expression for $\mathbb{E}\{(S_m^2)^2\}$ gives

$$\begin{aligned}\mathbb{E}\{(S_m^2)^2\} &= \frac{1}{(m-1)^2}\left[\left(1-\frac{2}{m}\right)\{m\mu_4 + m(m-1)v^2\} + \frac{1}{m^2}\{m\mu_4 + 3m(m-1)v^2\}\right] \\ &= \frac{\mu_4}{m} + \frac{m^2-2m+3}{m(m-1)}v^2.\end{aligned}$$

Since $\mathbb{E}(S_m^2) = v$, it then follows that

$$\begin{aligned}\mathrm{Var}(S_m^2) &= \mathbb{E}\{(S_m^2)^2\} - \{\mathbb{E}(S_m^2)\}^2 \\ &= \frac{\mu_4}{m} + \frac{m^2-2m+3}{m(m-1)}v^2 - v^2 \\ &= \frac{1}{m}\left(\mu_4 - \frac{m-3}{m-1}v^2\right).\end{aligned}$$

Therefore, the joint sampling covariance matrix is

$$\mathrm{Var}\begin{pmatrix}\bar{Y}_m \\ S_m^2\end{pmatrix} = \frac{1}{m}\begin{pmatrix} v & \mu_3 \\ \mu_3 & \mu_4 - \dfrac{m-3}{m-1}v^2\end{pmatrix},$$

which proves Equation (1). If $\mu_3 = 0$, the covariance between $\bar{Y}_m$ and S_m^2 vanishes. Under Gaussian sampling, the classical independence of the sample mean and sample variance gives the stronger conclusion that $\bar{Y}_m$ and S_m^2 are independent.

A.2 Proof of Proposition 2

Proof. Fix $\boldsymbol{x}$ and suppress it from the notation. The sample mean and unbiased sample variance have influence functions $\psi_\mu(Y) = Y - \mu$ and $\psi_v(Y) = (Y-\mu)^2 - v$. Write $\bar{X}_m = m^{-1}\sum_{j=1}^{m}(Y_j - \mu)$.

The unbiased sample variance satisfies

$$S_m^2 - v = \frac{m}{m-1}\left[\frac{1}{m}\sum_{j=1}^{m}\{(Y_j-\mu)^2 - v\} - \bar{X}_m^2\right] + \frac{v}{m-1}.$$

Under Assumption 1, $\bar{X}_m = O_p(m^{-1/2})$, and hence

$$\sqrt{m}(S_m^2 - v) = \frac{1}{\sqrt{m}}\sum_{j=1}^{m}\{(Y_j-\mu)^2 - v\} + o_p(1).$$

Together with

$$\sqrt{m}(\bar{Y}_m - \mu) = \frac{1}{\sqrt{m}}\sum_{j=1}^{m}(Y_j - \mu),$$

this gives the joint asymptotic linear representation

$$\sqrt{m}\begin{pmatrix}\bar{Y}_m - \mu \\ S_m^2 - v\end{pmatrix} = \frac{1}{\sqrt{m}}\sum_{j=1}^{m}\begin{pmatrix}\psi_\mu(Y_j) \\ \psi_v(Y_j)\end{pmatrix} + o_p(1).$$

The two influence functions have finite second moments. The multivariate central limit theorem therefore yields Equation (2), with

$$\mathbb{E}\begin{pmatrix}\psi_\mu^2 & \psi_\mu\psi_v \\ \psi_\mu\psi_v & \psi_v^2\end{pmatrix} = \begin{pmatrix}v & \mu_3 \\ \mu_3 & \mu_4 - v^2\end{pmatrix} = \boldsymbol{\Omega}_\infty.$$

Proposition 1 gives the exact finite-m covariance. Its lower-right entry differs from $(\mu_4 - v^2)/m$ by $2v^2/\{m(m-1)\}$, proving the expression for $\boldsymbol{\Omega}_m/m$ in the main text. Since $\mathbb{E}\bar{Y}_m = \mu$ and $\mathbb{E}S_m^2 = v$, Equation (3) matches the exact mean and covariance, while its Gaussian shape remains a working approximation. □

A.3 Sampling Covariance Estimation at Candidate Inputs

Evaluation of the acquisition function requires the sampling covariance of a future mean–variance estimate at any candidate $\boldsymbol{x} \in \mathcal{X}$. We first estimate the covariance structure from the simulator outputs already available at the evaluated design points and then extend these estimates to unevaluated candidate design points by inverse distance weighting (IDW). The covariance associated with a proposed replication count is obtained from the finite-replication correction in Proposition 1.

Suppose that the distinct evaluated design points are $\mathcal{H}_n = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n\}$. At $\boldsymbol{x}_i$, let $Y_{i1}, \ldots, Y_{im_i}$ denote the $m_i \geq 2$ available simulator outputs, with sample mean $\bar{Y}_i$ and unbiased sample variance S_i^2. Define

$$\widehat{\boldsymbol{\psi}}_{ij} = \begin{pmatrix}Y_{ij} - \bar{Y}_i \\ (Y_{ij} - \bar{Y}_i)^2 - S_i^2\end{pmatrix}, \qquad \widehat{\boldsymbol{\psi}}_{i\cdot} = \frac{1}{m_i}\sum_{j=1}^{m_i}\widehat{\boldsymbol{\psi}}_{ij}.$$

The estimator of the asymptotic sampling covariance at $\boldsymbol{x}_i$ is

$$\widehat{\boldsymbol{\Omega}}_{\infty,i} = \frac{1}{m_i - 1}\sum_{j=1}^{m_i}\left(\widehat{\boldsymbol{\psi}}_{ij} - \widehat{\boldsymbol{\psi}}_{i\cdot}\right)\left(\widehat{\boldsymbol{\psi}}_{ij} - \widehat{\boldsymbol{\psi}}_{i\cdot}\right)^\top.$$

Its population counterpart is

$$\boldsymbol{\Omega}_\infty(\boldsymbol{x}_i) = \begin{pmatrix}v(\boldsymbol{x}_i) & \mu_3(\boldsymbol{x}_i) \\ \mu_3(\boldsymbol{x}_i) & \mu_4(\boldsymbol{x}_i) - v^2(\boldsymbol{x}_i)\end{pmatrix}.$$

In particular, the off-diagonal entry estimates the third central moment $\mu_3(\boldsymbol{x}_i)$, which determines the mean–variance sampling covariance $\mu_3(\boldsymbol{x}_i)/m_i$.

For m_i available simulator outputs, the finite-replication estimator is

$$\widehat{\boldsymbol{\Omega}}_{m_i,i} = \widehat{\boldsymbol{\Omega}}_{\infty,i} + \begin{pmatrix}0 & 0 \\ 0 & \dfrac{2(S_i^2)^2}{m_i - 1}\end{pmatrix}.$$

Therefore, the estimated sampling covariance of $\boldsymbol{y}_i = (\bar{Y}_i, S_i^2)^\top$ is $\widehat{\boldsymbol{\Omega}}_{m_i,i}/m_i$. The following result shows that the plug-in estimator consistently recovers the finite-replication sampling covariance used in the observation model.

Lemma 1. *Under Assumption 1, for every fixed $\boldsymbol{x} \in \mathcal{X}$, as $m \to \infty$,*

$$\widehat{\boldsymbol{\Omega}}_m(\boldsymbol{x}) - \boldsymbol{\Omega}_m(\boldsymbol{x}) \xrightarrow{p} 0, \qquad \frac{\widehat{\boldsymbol{\Omega}}_m(\boldsymbol{x})}{m} = \frac{\boldsymbol{\Omega}_m(\boldsymbol{x})}{m} + o_p(m^{-1}).$$

Proof. Fix $\boldsymbol{x}$ and suppress it from the notation. Let $e_j = Y_j - \bar{Y}_m$ and define the empirical central moments $\widehat{\mu}_{k,m} = m^{-1}\sum_{j=1}^m e_j^k$ for $k = 2, 3, 4$. The estimator of the asymptotic sampling covariance can then be written as

$$\widehat{\boldsymbol{\Omega}}_\infty = \frac{m}{m-1}\begin{pmatrix} \widehat{\mu}_{2,m} & \widehat{\mu}_{3,m} \\ \widehat{\mu}_{3,m} & \widehat{\mu}_{4,m} - \widehat{\mu}_{2,m}^2 \end{pmatrix}.$$

Under Assumption 1, $\widehat{\mu}_{2,m} \xrightarrow{p} v$, $\widehat{\mu}_{3,m} \xrightarrow{p} \mu_3$, and $\widehat{\mu}_{4,m} \xrightarrow{p} \mu_4$. Since $m/(m-1) \to 1$, it follows that $\widehat{\boldsymbol{\Omega}}_\infty - \boldsymbol{\Omega}_\infty \xrightarrow{p} 0$.

Moreover, $S_m^2 \to_p v$, and hence

$$\frac{2(S_m^2)^2}{m-1} - \frac{2v^2}{m-1} = o_p(1).$$

Combining the two results gives $\widehat{\boldsymbol{\Omega}}_m - \boldsymbol{\Omega}_m = o_p(1)$. Dividing by m proves the second assertion. □

We next extend the design-point-specific estimates to unevaluated candidate design points. All distances are computed in the scaled input space used by the GP. Let $p_{\mathrm{idw}} > 0$ denote the IDW power parameter, which controls the rate at which the weight assigned to an evaluated design point decreases with distance. For $\boldsymbol{x} \notin \mathcal{H}_n$, define $d_i(\boldsymbol{x}) = \|\boldsymbol{x} - \boldsymbol{x}_i\|_2$ and

$$w_i(\boldsymbol{x}) = \frac{d_i(\boldsymbol{x})^{-p_{\mathrm{idw}}}}{\sum_{k=1}^n d_k(\boldsymbol{x})^{-p_{\mathrm{idw}}}}, \qquad i = 1, \ldots, n.$$

The IDW estimates are

$$\widehat{\boldsymbol{\Omega}}_{\infty,n}(\boldsymbol{x}) = \sum_{i=1}^n w_i(\boldsymbol{x})\widehat{\boldsymbol{\Omega}}_{\infty,i}, \qquad \widehat{v}_n(\boldsymbol{x}) = \sum_{i=1}^n w_i(\boldsymbol{x}) S_i^2, \qquad \boldsymbol{x} \notin \mathcal{H}_n.$$

At an evaluated design point $\boldsymbol{x}_i$, the corresponding design-point-specific estimates are used directly: $\widehat{\boldsymbol{\Omega}}_{\infty,n}(\boldsymbol{x}_i) = \widehat{\boldsymbol{\Omega}}_{\infty,i}$ and $\widehat{v}_n(\boldsymbol{x}_i) = S_i^2$.

For a future evaluation with $r \geq 2$ replications at $\boldsymbol{x}$, define

$$\widehat{\boldsymbol{\Omega}}_{n,r}(\boldsymbol{x}) = \widehat{\boldsymbol{\Omega}}_{\infty,n}(\boldsymbol{x}) + \begin{pmatrix} 0 & 0 \\ 0 & \dfrac{2\widehat{v}_n^2(\boldsymbol{x})}{r-1} \end{pmatrix}.$$

The estimated sampling covariance of the future mean–variance estimate is $\widehat{\boldsymbol{\Omega}}_{n,r}(\boldsymbol{x})/r$. This construction combines the estimated asymptotic covariance at $\boldsymbol{x}$ with the finite-replication correction associated with the proposed value of r.

Finally, the sampling covariance is transformed to the scale of the standardized latent objectives. Recall that $\boldsymbol{f}_s(\boldsymbol{x}) = \mathbf{W}\{\boldsymbol{f}(\boldsymbol{x}) - \boldsymbol{c}_0\}$, where $\mathbf{W}$ is a fixed positive diagonal matrix. If the initial design contains n_0 design points with m_0 replications per design point, define

$$\widehat{\tau}_0^2 = \frac{1}{2n_0}\sum_{i=1}^{n_0} \mathrm{tr}\left(\mathbf{W}\widehat{\boldsymbol{\Omega}}_{m_0,i}\mathbf{W}^\top\right).$$

For a candidate $\boldsymbol{x}$ and replication count r, the standardized relative scaled sampling covariance matrix is

$$\widehat{\mathbf{Q}}_{n,r}(\boldsymbol{x}) = \frac{\mathbf{W}\widehat{\boldsymbol{\Omega}}_{n,r}(\boldsymbol{x})\mathbf{W}^\top}{\widehat{\tau}_0^2}.$$

The conditional covariance of the standardized mean–variance estimate is $(\sigma^2\eta/r)\widehat{\mathbf{Q}}_{n,r}(\boldsymbol{x})$, as specified in the hierarchical GP model. The parameter η determines the magnitude of the sampling covariance relative to the latent-process covariance.

Since each $\widehat{\boldsymbol{\Omega}}_{\infty,i}$ is positive semidefinite, $w_i(\boldsymbol{x}) \geq 0$, and $\sum_{i=1}^n w_i(\boldsymbol{x}) = 1$, it follows that $\widehat{\boldsymbol{\Omega}}_{n,r}(\boldsymbol{x}) \succeq 0$ and $\widehat{\mathbf{Q}}_{n,r}(\boldsymbol{x}) \succeq 0$. For small replication counts these estimates can be singular; when strict positive definiteness is required for numerical inversion, we add a fixed diagonal ridge $\epsilon_\Omega I_2$, where $\epsilon_\Omega > 0$ is held constant throughout a run.

B Hierarchical Posterior

We impose the normalization $\mathrm{tr}\{\mathbf{K}_\vartheta(\boldsymbol{x},\boldsymbol{x})\} = 2$ on the stationary kernel. Under this convention, σ^2 controls the overall scale of the latent process, whereas η controls the sampling covariance relative to that scale.

B.1 Covariance Fitting

Conditional on the estimated sampling-covariance matrices, write $\mathbf{C}_n = \mathbf{C}_n(\vartheta,\eta)$. From Equation (4), the log-likelihood is

$$\ell(\boldsymbol{\beta},\sigma^2,\vartheta,\eta) = -\frac{N}{2}\log(2\pi) - \frac{N}{2}\log\sigma^2 - \frac{1}{2}\log|\mathbf{C}_n| - \frac{(\boldsymbol{y}_s - \mathbf{H}\boldsymbol{\beta})^\top \mathbf{C}_n^{-1}(\boldsymbol{y}_s - \mathbf{H}\boldsymbol{\beta})}{2\sigma^2}.$$

For fixed (ϑ,η), differentiation with respect to $\boldsymbol{\beta}$ gives $\mathbf{H}^\top\mathbf{C}_n^{-1}(\boldsymbol{y}_s - \mathbf{H}\widehat{\boldsymbol{\beta}}_n) = 0$. Therefore,

$$\widehat{\boldsymbol{\beta}}_n = (\mathbf{H}^\top\mathbf{C}_n^{-1}\mathbf{H})^{-1}\mathbf{H}^\top\mathbf{C}_n^{-1}\boldsymbol{y}_s.$$

Let $R_n(\vartheta,\eta) = (\boldsymbol{y}_s - \mathbf{H}\widehat{\boldsymbol{\beta}}_n)^\top\mathbf{C}_n^{-1}(\boldsymbol{y}_s - \mathbf{H}\widehat{\boldsymbol{\beta}}_n)$. Differentiation with respect to σ^2 then gives $\widehat{\sigma}_n^2 = R_n(\vartheta,\eta)/N$. Substitution into the log-likelihood yields

$$-2\ell_{\mathrm{prof}}(\vartheta,\eta) = N\log\widehat{\sigma}_n^2 + \log|\mathbf{C}_n| + N\{1+\log(2\pi)\}.$$

It follows that

$$(\widehat{\vartheta}_n,\widehat{\eta}_n) \in \arg\min_{\vartheta,\eta}\left\{N\log\widehat{\sigma}_n^2 + \log|\mathbf{C}_n|\right\},$$

where $\eta > 0$ and $\mathbf{C}_n$ is required to be positive definite. The fitted values $(\widehat{\vartheta}_n,\widehat{\eta}_n)$ are held fixed in the posterior calculations and are re-estimated after each new evaluation.

B.2 Predictive Quantities

All covariance quantities below are evaluated at the fitted parameters. Define $\mathbf{K}_{\boldsymbol{x},n} = [\mathbf{K}_{\widehat{\vartheta}_n}(\boldsymbol{x},\boldsymbol{x}_1),\ldots,\mathbf{K}_{\widehat{\vartheta}_n}(\boldsymbol{x},\boldsymbol{x}_n)]$, $\mathbf{A}_n = \mathbf{H}^\top\mathbf{C}_n^{-1}\mathbf{H}$, and $\boldsymbol{\Delta}_n(\boldsymbol{x}) = \mathbf{H}(\boldsymbol{x}) - \mathbf{K}_{\boldsymbol{x},n}\mathbf{C}_n^{-1}\mathbf{H}$. Also define $\mathbf{V}_{s,n}(\boldsymbol{x}) = \mathbf{K}_{\widehat{\vartheta}_n}(\boldsymbol{x},\boldsymbol{x}) - \mathbf{K}_{\boldsymbol{x},n}\mathbf{C}_n^{-1}\mathbf{K}_{\boldsymbol{x},n}^\top$. The predictive location and covariance factor on the standardized scale are

$$\boldsymbol{m}_{s,n}(\boldsymbol{x}) = \mathbf{H}(\boldsymbol{x})\widehat{\boldsymbol{\beta}}_n + \mathbf{K}_{\boldsymbol{x},n}\mathbf{C}_n^{-1}(\boldsymbol{y}_s - \mathbf{H}\widehat{\boldsymbol{\beta}}_n)$$

and

$$\mathbf{S}_{s,n}(\boldsymbol{x}) = \mathbf{V}_{s,n}(\boldsymbol{x}) + \boldsymbol{\Delta}_n(\boldsymbol{x})\mathbf{A}_n^{-1}\boldsymbol{\Delta}_n(\boldsymbol{x})^\top.$$

Here, $\mathbf{V}_{s,n}(\boldsymbol{x})$ is the conditional predictive scale matrix for a fixed trend coefficient, and the second term accounts for posterior uncertainty in $\boldsymbol{\beta}$.

Using the fixed transformation from Appendix A.3, define $\boldsymbol{m}_n(\boldsymbol{x}) = \boldsymbol{c}_0 + \mathbf{W}^{-1}\boldsymbol{m}_{s,n}(\boldsymbol{x})$ and $\mathbf{S}_n(\boldsymbol{x}) = \mathbf{W}^{-1}\mathbf{S}_{s,n}(\boldsymbol{x})\mathbf{W}^{-\top}$. Conditional on σ^2, the predictive distribution on the original objective scale is $\boldsymbol{f}(\boldsymbol{x}) \mid \sigma^2, \mathcal{D}_n \sim \mathcal{N}_2\{\boldsymbol{m}_n(\boldsymbol{x}), \sigma^2\mathbf{S}_n(\boldsymbol{x})\}$.

The scale matrix of the marginal Student-t distribution is $\widetilde{\sigma}_n^2\mathbf{S}_n(\boldsymbol{x})$. If $\nu_n > 2$, its covariance matrix is $\nu_n\widetilde{\sigma}_n^2\mathbf{S}_n(\boldsymbol{x})/(\nu_n - 2)$.

B.3 Proof of Proposition 3

Throughout the proof, the fitted covariance parameters and estimated sampling-covariance matrices are treated as fixed. By Equation (4),

$$p(\boldsymbol{y}_s \mid \boldsymbol{\beta}, \sigma^2) \propto (\sigma^2)^{-N/2} |\mathbf{C}_n|^{-1/2} \exp\left\{ -\frac{(\boldsymbol{y}_s - \mathbf{H}\boldsymbol{\beta})^\top \mathbf{C}_n^{-1} (\boldsymbol{y}_s - \mathbf{H}\boldsymbol{\beta})}{2\sigma^2} \right\}.$$

The generalized least-squares normal equation implies $\mathbf{H}^\top \mathbf{C}_n^{-1}(\boldsymbol{y}_s - \mathbf{H}\widehat{\boldsymbol{\beta}}_n) = 0$. Hence,

$$(\boldsymbol{y}_s - \mathbf{H}\boldsymbol{\beta})^\top \mathbf{C}_n^{-1} (\boldsymbol{y}_s - \mathbf{H}\boldsymbol{\beta}) = R_n + (\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}_n)^\top \mathbf{A}_n (\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}_n).$$

Combining the likelihood with $\pi(\boldsymbol{\beta}) \propto 1$ and $p(\sigma^2) \propto (\sigma^2)^{-a-1} \exp(-b/\sigma^2)$ gives

$$p(\boldsymbol{\beta}, \sigma^2 \mid \mathcal{D}_n) \propto (\sigma^2)^{-(N/2+a+1)} \exp\left[-\frac{R_n + (\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}_n)^\top \mathbf{A}_n (\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}_n) + 2b}{2\sigma^2} \right].$$

Regarding this expression as a function of $\boldsymbol{\beta}$ gives $\boldsymbol{\beta} \mid \sigma^2, \mathcal{D}_n \sim \mathcal{N}\left(\widehat{\boldsymbol{\beta}}_n, \sigma^2 \mathbf{A}_n^{-1}\right)$.

To obtain the marginal posterior of σ^2, integrate out $\boldsymbol{\beta}$. Since

$$\int_{\mathbb{R}^{p_\beta}} \exp\left\{ -\frac{(\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}_n)^\top \mathbf{A}_n (\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}_n)}{2\sigma^2} \right\} d\boldsymbol{\beta} = (2\pi\sigma^2)^{p_\beta/2} |\mathbf{A}_n|^{-1/2},$$

it follows that

$$p(\sigma^2 \mid \mathcal{D}_n) \propto (\sigma^2)^{-a_n - 1} \exp\left(-\frac{b_n}{\sigma^2} \right),$$

where $a_n = a + (N - p_\beta)/2$ and $b_n = b + R_n/2$. Therefore, $\sigma^2 \mid \mathcal{D}_n \sim \mathrm{IG}(a_n, b_n)$.

We next derive the predictive distribution. Conditional on $(\boldsymbol{\beta}, \sigma^2)$, the joint distribution on the standardized scale is

$$\begin{pmatrix} \boldsymbol{y}_s \\ \boldsymbol{f}_s(\boldsymbol{x}) \end{pmatrix} \,\middle|\, \boldsymbol{\beta}, \sigma^2 \sim \mathcal{N}\left\{ \begin{pmatrix} \mathbf{H}\boldsymbol{\beta} \\ \mathbf{H}(\boldsymbol{x})\boldsymbol{\beta} \end{pmatrix}, \sigma^2 \begin{pmatrix} \mathbf{C}_n & \mathbf{K}_{\boldsymbol{x},n}^\top \\ \mathbf{K}_{\boldsymbol{x},n} & \mathbf{K}_{\widehat{\vartheta}_n}(\boldsymbol{x}, \boldsymbol{x}) \end{pmatrix} \right\}.$$

The conditional normal formula gives

$$\boldsymbol{f}_s(\boldsymbol{x}) \mid \boldsymbol{y}_s, \boldsymbol{\beta}, \sigma^2 \sim \mathcal{N}_2\left\{ \mathbf{K}_{\boldsymbol{x},n} \mathbf{C}_n^{-1} \boldsymbol{y}_s + \boldsymbol{\Delta}_n(\boldsymbol{x})\boldsymbol{\beta},\ \sigma^2 \mathbf{V}_{s,n}(\boldsymbol{x}) \right\}.$$

Since $\boldsymbol{\beta} \mid \sigma^2, \mathcal{D}_n \sim \mathcal{N}(\widehat{\boldsymbol{\beta}}_n, \sigma^2 \mathbf{A}_n^{-1})$, integration with respect to $\boldsymbol{\beta}$ yields

$$\boldsymbol{f}_s(\boldsymbol{x}) \mid \sigma^2, \mathcal{D}_n \sim \mathcal{N}_2\left\{ \boldsymbol{m}_{s,n}(\boldsymbol{x}), \sigma^2 \mathbf{S}_{s,n}(\boldsymbol{x}) \right\}.$$

It then follows that

$$\boldsymbol{f}(\boldsymbol{x}) \mid \sigma^2, \mathcal{D}_n \sim \mathcal{N}_2\left\{ \boldsymbol{m}_n(\boldsymbol{x}), \sigma^2 \mathbf{S}_n(\boldsymbol{x}) \right\}.$$

Finally, integrate σ^2 with respect to its inverse-Gamma posterior:

$$\begin{aligned} p\{\boldsymbol{f}(\boldsymbol{x}) \mid \mathcal{D}_n\} &= \int_0^\infty p\{\boldsymbol{f}(\boldsymbol{x}) \mid \sigma^2, \mathcal{D}_n\} p(\sigma^2 \mid \mathcal{D}_n)\, d\sigma^2 \\ &= \frac{\Gamma(a_n + 1)}{\Gamma(a_n)\, 2\pi b_n\, |\mathbf{S}_n(\boldsymbol{x})|^{1/2}} \left[1 + \frac{\{\boldsymbol{f}(\boldsymbol{x}) - \boldsymbol{m}_n(\boldsymbol{x})\}^\top \mathbf{S}_n(\boldsymbol{x})^{-1} \{\boldsymbol{f}(\boldsymbol{x}) - \boldsymbol{m}_n(\boldsymbol{x})\}}{2b_n} \right]^{-(a_n+1)}. \end{aligned}$$

This is the density of the bivariate non-standardized Student-t distribution

$$\boldsymbol{f}(\boldsymbol{x}) \mid \mathcal{D}_n \sim T_2\left\{ 2a_n, \boldsymbol{m}_n(\boldsymbol{x}), \frac{b_n}{a_n} \mathbf{S}_n(\boldsymbol{x}) \right\}.$$

Since $\nu_n = 2a_n$ and $\widetilde{\sigma}_n^2 = b_n/a_n$, the result follows from Equation (5).

C Prospective Updates and Reevaluation

C.1 Proof of Proposition 4

Condition on the current data $\mathcal{D}_n$ and fitted covariance quantities. Fix $(\boldsymbol{x}, r)$ and write $\boldsymbol{m} = \boldsymbol{m}_n(\boldsymbol{x})$, $\mathbf{S} = \mathbf{S}_n(\boldsymbol{x})$, $\boldsymbol{\Lambda} = \boldsymbol{\Lambda}_r(\boldsymbol{x})$, and $\widetilde{\boldsymbol{y}} = \widetilde{\boldsymbol{y}}(\boldsymbol{x}, r)$. Assume $\mathbf{S}, \boldsymbol{\Lambda} \succeq 0$ and $\mathbf{S} + \boldsymbol{\Lambda} \succ 0$. With $\mathbf{D}_0 = \mathbf{W}^{-1}$ from Appendix A.3, the observation scaled sampling covariance matrix on the original scale is $\boldsymbol{\Lambda}_r(\boldsymbol{x}) = (\widehat{\eta}_n/r)\mathbf{D}_0\widehat{\mathbf{Q}}_{n,r}(\boldsymbol{x})\mathbf{D}_0^\top$.

Under the working model, the sampling error in the future mean–variance estimate is conditionally independent of $\boldsymbol{f}(\boldsymbol{x})$ given σ^2, so

$$\begin{pmatrix} \boldsymbol{f}(\boldsymbol{x}) \\ \widetilde{\boldsymbol{y}} \end{pmatrix} \Big| \, \mathcal{D}_n, \sigma^2 \sim \mathcal{N}_4\left\{ \begin{pmatrix} \boldsymbol{m} \\ \boldsymbol{m} \end{pmatrix}, \sigma^2 \begin{pmatrix} \mathbf{S} & \mathbf{S} \\ \mathbf{S} & \mathbf{S} + \boldsymbol{\Lambda} \end{pmatrix} \right\}.$$

Gaussian conditioning gives

$$\begin{aligned} \boldsymbol{m}_+(\boldsymbol{x}, r) &= \boldsymbol{m} + \mathbf{S}(\mathbf{S} + \boldsymbol{\Lambda})^{-1}(\widetilde{\boldsymbol{y}} - \boldsymbol{m}), \\ \mathbf{S}_+(\boldsymbol{x}, r) &= \mathbf{S} - \mathbf{S}(\mathbf{S} + \boldsymbol{\Lambda})^{-1}\mathbf{S}. \end{aligned}$$

Since $\boldsymbol{m}_+(\boldsymbol{x}, r)$ is affine in $\widetilde{\boldsymbol{y}} \mid \mathcal{D}_n, \sigma^2 \sim \mathcal{N}_2\{\boldsymbol{m}, \sigma^2(\mathbf{S} + \boldsymbol{\Lambda})\}$, it follows that $\boldsymbol{m}_+(\boldsymbol{x}, r) \mid \mathcal{D}_n, \sigma^2 \sim \mathcal{N}_2\{\boldsymbol{m}, \sigma^2\mathbf{U}_r(\boldsymbol{x})\}$, where $\mathbf{U}_r(\boldsymbol{x}) = \mathbf{S}(\mathbf{S} + \boldsymbol{\Lambda})^{-1}\mathbf{S} = \mathbf{S} - \mathbf{S}_+(\boldsymbol{x}, r)$. Integrating $\sigma^2 \mid \mathcal{D}_n \sim \mathrm{IG}(a_n, b_n)$ gives $\boldsymbol{m}_+(\boldsymbol{x}, r) \mid \mathcal{D}_n \sim T_2\{\nu_n, \boldsymbol{m}, \widetilde{\sigma}_n^2\mathbf{U}_r(\boldsymbol{x})\}$, with $\nu_n = 2a_n$ and $\widetilde{\sigma}_n^2 = b_n/a_n$. The scale-mixture interpretation also covers singular $\mathbf{U}_r(\boldsymbol{x})$.

With ν_n and $\widetilde{\sigma}_n$ fixed, the displacement $\boldsymbol{\Xi}_+(\boldsymbol{x}, r) - \boldsymbol{m}_+(\boldsymbol{x}, r)$ is deterministic. Translation therefore gives Equation (7). Equation (6) follows directly from the univariate Student-t margins.

The updated quantile vector integrates over the future observation conditional on the currently fitted covariance parameters. When a previously evaluated design point is reevaluated, the new set of replications is a working observation conditionally independent of previous simulator outputs given $\boldsymbol{f}(\boldsymbol{x})$ and σ^2. After the simulator outputs are obtained, all outputs at that design point are pooled and the surrogate is refitted, so the acquisition calculation need not coincide with the refitted model.

C.2 Pooling Raw Replications after Reevaluation

Let $Y_1, \ldots, Y_m$ be the available simulator outputs and $Z_1, \ldots, Z_r$ a new set of replications, with $m, r \geq 2$. Write $\bar{Y}_m$ and $\bar{Z}_r$ for their sample means, $M_{2,m} = \sum_{j=1}^m (Y_j - \bar{Y}_m)^2$, and $M_{2,r}^{\mathrm{new}} = \sum_{j=1}^r (Z_j - \bar{Z}_r)^2$. The pooled mean is $\bar{Y}_{m+r} = (m\bar{Y}_m + r\bar{Z}_r)/(m + r)$. Expanding the squared deviations about the two sample means and using their zero residual sums gives

$$M_{2,m+r} = M_{2,m} + M_{2,r}^{\mathrm{new}} + \frac{mr}{m+r}(\bar{Y}_m - \bar{Z}_r)^2.$$

Hence $S_{m+r}^2 = M_{2,m+r}/(m + r - 1)$. These identities hold samplewise without a Gaussian assumption. The pooled raw outputs are also used to recompute the third and fourth empirical moments and the design-point-specific covariance estimate. The new set of replications is therefore not treated as a separate design-point-level observation after refitting.

C.3 Hypervolume Improvement under Reevaluation

Let $\mathcal{A}_n = \{\boldsymbol{\xi}_n(\boldsymbol{x}_k) : k = 1, \ldots, n\}$, with $\mathcal{P}_n^{(\tau)} = \mathrm{Pareto}(\mathcal{A}_n)$. For reevaluation of $\boldsymbol{x}_i$, define $\mathcal{A}_{n,-i} = \{\boldsymbol{\xi}_n(\boldsymbol{x}_k) : k \neq i\}$ and $\mathcal{C}_n(\boldsymbol{x}_i) = \mathrm{Pareto}(\mathcal{A}_{n,-i})$. Removal is by design-point index, so an identical vector at another design point remains. The removed contribution is $\delta_n(\boldsymbol{x}_i) = \mathrm{HV}\{\mathcal{P}_n^{(\tau)}; \boldsymbol{r}_{\mathrm{ref}}\} - \mathrm{HV}\{\mathcal{C}_n(\boldsymbol{x}_i); \boldsymbol{r}_{\mathrm{ref}}\} \geq 0$. For a new design point, set $\mathcal{C}_n(\boldsymbol{x}) = \mathcal{P}_n^{(\tau)}$ and $\delta_n(\boldsymbol{x}) = 0$. Recall that $\mathrm{HVI}(\boldsymbol{z}; \mathcal{C}, \boldsymbol{r}_{\mathrm{ref}}) = \mathrm{HV}(\mathcal{C} \cup \{\boldsymbol{z}\}; \boldsymbol{r}_{\mathrm{ref}}) - \mathrm{HV}(\mathcal{C}; \boldsymbol{r}_{\mathrm{ref}})$. Removing dominated vectors leaves hypervolume unchanged. Thus, for every updated quantile vector $\boldsymbol{z}$,

$$\begin{aligned} &\mathrm{HV}\{\mathcal{P}_+(\boldsymbol{x}_i, \boldsymbol{z}); \boldsymbol{r}_{\mathrm{ref}}\} - \mathrm{HV}\{\mathcal{P}_n^{(\tau)}; \boldsymbol{r}_{\mathrm{ref}}\} \\ &\quad = \mathrm{HV}\{\mathcal{C}_n(\boldsymbol{x}_i) \cup \{\boldsymbol{z}\}; \boldsymbol{r}_{\mathrm{ref}}\} - \mathrm{HV}\{\mathcal{C}_n(\boldsymbol{x}_i); \boldsymbol{r}_{\mathrm{ref}}\} - \delta_n(\boldsymbol{x}_i) \\ &\quad = \mathrm{HVI}(\boldsymbol{z}; \mathcal{C}_n(\boldsymbol{x}_i), \boldsymbol{r}_{\mathrm{ref}}) - \delta_n(\boldsymbol{x}_i). \end{aligned}$$

The same identity holds at a new design point with $\delta_n(\boldsymbol{x}) = 0$. Substituting $\boldsymbol{\Xi}_+(\boldsymbol{x}, r)$ and taking the conditional expectation of the positive part gives the equivalent HVI representation of Equation (8). Moreover, $0 \le [\mathrm{HVI}(\boldsymbol{z}; \mathcal{C}_n(\boldsymbol{x}), \boldsymbol{r}_{\mathrm{ref}}) - \delta_n(\boldsymbol{x})]_+ \le [r_{\mathrm{ref},1} - \boldsymbol{z}_1]_+ [r_{\mathrm{ref},2} - \boldsymbol{z}_2]_+$. Since $ab \le (a^2 + b^2)/2$, the finite second moments of Equation (7) ensure integrability when $\nu_n > 2$. The remaining set must include all design points, since a previously dominated vector may become nondominated after removal. This calculation holds the other quantile vectors fixed; it does not describe their simultaneous changes after refitting.

C.4 Posterior Covariance under Additional Replications

The next result formalizes how increasing the replication count affects the observation covariance and the corresponding posterior update under the currently fitted model.

Lemma 2. *Fix $\boldsymbol{x}$ and the fitted covariance quantities, with $\mathbf{S}_n(\boldsymbol{x})$ and $\boldsymbol{\Lambda}_r(\boldsymbol{x})$ positive definite for all candidate replication counts. For fixed plug-in moments and normalization, Appendix A.3 gives $\widehat{\mathbf{Q}}_{n,r}(\boldsymbol{x}) = \mathbf{Q}_{\infty,n}(\boldsymbol{x}) + \mathbf{B}_n(\boldsymbol{x})/(r-1)$, where $\mathbf{Q}_{\infty,n}(\boldsymbol{x}) \succeq 0$ and $\mathbf{B}_n(\boldsymbol{x}) \succeq 0$ are the normalized asymptotic scaled sampling covariance matrix and finite-replication correction, respectively. For $2 \le r_1 \le r_2$,*

$$\boldsymbol{\Lambda}_{r_2}(\boldsymbol{x}) \preceq \boldsymbol{\Lambda}_{r_1}(\boldsymbol{x}), \qquad \mathbf{S}_+(\boldsymbol{x}, r_2) \preceq \mathbf{S}_+(\boldsymbol{x}, r_1), \qquad U_{r_2}(\boldsymbol{x}) \succeq U_{r_1}(\boldsymbol{x}),$$

where $\preceq$ denotes the Loewner order.

Proof. The decomposition gives $\boldsymbol{\Lambda}_r(\boldsymbol{x}) = \widehat{\eta}_n \mathbf{D}_0 \{\mathbf{Q}_{\infty,n}(\boldsymbol{x})/r + \mathbf{B}_n(\boldsymbol{x})/[r(r-1)]\} \mathbf{D}_0^\top$. Both scalar coefficients decrease with $r \ge 2$, and congruence preserves positive semidefiniteness, proving the first ordering. Since matrix inversion reverses the Loewner order, $\mathbf{S}_+(\boldsymbol{x}, r) = \{\mathbf{S}_n(\boldsymbol{x})^{-1} + \boldsymbol{\Lambda}_r(\boldsymbol{x})^{-1}\}^{-1}$ gives the second. The third follows from $\mathbf{U}_r(\boldsymbol{x}) = \mathbf{S}_n(\boldsymbol{x}) - \mathbf{S}_+(\boldsymbol{x}, r)$. □

These orderings concern conditional covariance under the same fitted model; they do not establish monotonicity of $\mathrm{HEQHVI}_n(\boldsymbol{x}, r)/r$.

D Numerical Evaluation and Optimization of the Acquisition

D.1 Serial Integration and Pathwise Differentiation

For a fixed pair $(\boldsymbol{x}, r)$, let $\boldsymbol{\ell}_n(\boldsymbol{x}, r) = \boldsymbol{m}_n(\boldsymbol{x}) + c_{\tau,n} \widetilde{\sigma}_n (\sqrt{[\mathbf{S}_+(\boldsymbol{x}, r)]_{11}}, \sqrt{[\mathbf{S}_+(\boldsymbol{x}, r)]_{22}})^\top$ and $\mathbf{V}_n(\boldsymbol{x}, r) = \widetilde{\sigma}_n^2 \mathbf{U}_r(\boldsymbol{x})$, as in Equation (7). Choose $\mathbf{L}_n \mathbf{L}_n^\top = \mathbf{V}_n$. With common draws $\boldsymbol{z}_s \sim \mathcal{N}_2(0, I_2)$ and $g_s \sim \chi^2_{\nu_n}$, define $\boldsymbol{e}_s = \boldsymbol{z}_s \sqrt{\nu_n / g_s}$ and $\boldsymbol{\Xi}_+^{(s)} = \boldsymbol{\ell}_n + \mathbf{L}_n \boldsymbol{e}_s$. Hence $\boldsymbol{\Xi}_+^{(s)}$ has the required bivariate Student law.

Write $G_n(\boldsymbol{\xi}; \boldsymbol{x}) = [\mathrm{HVI}\{\boldsymbol{\xi}; \mathcal{C}_n(\boldsymbol{x}), \boldsymbol{r}_{\mathrm{ref}}\} - \delta_n(\boldsymbol{x})]_+$. Let $c(\boldsymbol{x}, r)$ denote the evaluation cost; in this work, $c(\boldsymbol{x}, r) = r$. Conditional on the common draws, the cost-adjusted sample-average approximation is

$$\widehat{\alpha}_{n,M}(\boldsymbol{x}, r) = \{M c(\boldsymbol{x}, r)\}^{-1} \sum_{s=1}^{M} G_n\{\boldsymbol{\ell}_n(\boldsymbol{x}, r) + \mathbf{L}_n(\boldsymbol{x}, r) \boldsymbol{e}_s; \boldsymbol{x}\}.$$

The same draws are kept fixed throughout a BO iteration, so this is a deterministic function of $(\boldsymbol{x}, r)$ during acquisition optimization.

For a new design point, $\mathcal{C}_n(\boldsymbol{x}) = \mathcal{P}_n^{(\tau)}$ and $\delta_n(\boldsymbol{x}) = 0$. Let $J_s(\boldsymbol{x}, r) = \partial \boldsymbol{\Xi}_+^{(s)}(\boldsymbol{x}, r) / \partial \boldsymbol{x}^\top$. Away from Pareto-cell and positive-part boundaries, the chain rule gives

$$\nabla_{\boldsymbol{x}} \widehat{\alpha}_{n,M}(\boldsymbol{x}, r) = \frac{1}{M c(\boldsymbol{x}, r)} \sum_{s=1}^{M} J_s(\boldsymbol{x}, r)^\top \nabla_{\boldsymbol{\xi}} G_n\{\boldsymbol{\Xi}_+^{(s)}(\boldsymbol{x}, r); \boldsymbol{x}\} - \widehat{\alpha}_{n,M}(\boldsymbol{x}, r) \nabla_{\boldsymbol{x}} \log c(\boldsymbol{x}, r), \tag{14}$$

where

$$J_s(:, k) = \frac{\partial \boldsymbol{\ell}_n}{\partial \boldsymbol{x}_k} + \frac{\partial \mathbf{L}_n}{\partial \boldsymbol{x}_k} \boldsymbol{e}_s, \qquad \nabla_{\boldsymbol{\xi}} G_n = \mathbf{1}\{\mathrm{HVI} > \delta_n\} \nabla_{\boldsymbol{\xi}} \mathrm{HVI}\,.$$

At points where the positive-part operator is nondifferentiable, we use a measurable subgradient. The derivatives of $\boldsymbol{\ell}_n$ and the 2×2 factor $\mathbf{L}_n$ are evaluated by centered differences, with one-sided differences on the boundary. If $c(\boldsymbol{x}, r) = r$, the last term in Equation (14) is zero. Previously evaluated design points are evaluated by the discrete reevaluation branch and do not enter this continuous derivative.

For each $r \in \mathcal{R}$, constrained optimization of the sample-average criterion produces a candidate set. Final scores use normalized generalized Gauss–Laguerre weights w_a^{Γ} for the scale mixture and tensor Gauss–Hermite weights $w_{b_1}^H w_{b_2}^H$ for the Gaussian coordinates:

$$\widehat{\alpha}_n^{\mathrm{quad}}(\boldsymbol{x}, r) = \frac{1}{c(\boldsymbol{x}, r)} \sum_{a, b_1, b_2} w_a^{\Gamma} w_{b_1}^H w_{b_2}^H G_n\{\boldsymbol{\ell}_n(\boldsymbol{x}, r) + \mathbf{L}_n(\boldsymbol{x}, r) \boldsymbol{e}_{a b_1 b_2}; \boldsymbol{x}\}.$$

The maximizer of this common quadrature score over the candidate set, eligible reevaluations, and $r \in \mathcal{R}$ is selected. The serial policy and the primary pair of the parallel policy use identical nodes, weights, and tie rules.

D.2 Parallel MED Search

For batch position $j \geq 2$, enumerate the positive terms in $\mathcal{L}_{n,t,j}(\boldsymbol{x}) = \{\|\boldsymbol{x} - \boldsymbol{z}\|_2 : \boldsymbol{z} \in \mathcal{H}_n\} \cup \{\bar{\alpha}_n(\tilde{\boldsymbol{x}}_\ell)\|\boldsymbol{x} - \tilde{\boldsymbol{x}}_\ell\|_2 : \ell < j\}$ as $h_1(\boldsymbol{x}), \ldots, h_K(\boldsymbol{x})$. The power soft minimum and its normalized weights are

$$D_\rho(\boldsymbol{x}) = \left\{ \sum_{k=1}^K h_k(\boldsymbol{x})^{-\rho} \right\}^{-1/\rho}, \qquad \pi_k(\boldsymbol{x}) = \frac{h_k(\boldsymbol{x})^{-\rho}}{\sum_{l=1}^K h_l(\boldsymbol{x})^{-\rho}}.$$

Since $\log D_\rho = -\rho^{-1} \log \sum_k h_k^{-\rho}$, logarithmic differentiation yields

$$\nabla_x D_\rho(\boldsymbol{x}) = D_\rho(\boldsymbol{x}) \sum_{k=1}^K \pi_k(\boldsymbol{x}) \nabla_{\boldsymbol{x}} \log h_k(\boldsymbol{x}), \qquad D_\rho(\boldsymbol{x}) \longrightarrow \min_k h_k(\boldsymbol{x}) \quad (\rho \to \infty).$$

In particular, $\nabla_{\boldsymbol{x}} \|\boldsymbol{x} - \boldsymbol{z}\|_2 = (\boldsymbol{x} - \boldsymbol{z})/\|\boldsymbol{x} - \boldsymbol{z}\|_2$, while a within-batch term has gradient $\bar{\alpha}_n(\tilde{\boldsymbol{x}}_\ell)(\boldsymbol{x} - \tilde{\boldsymbol{x}}_\ell)/\|\boldsymbol{x} - \tilde{\boldsymbol{x}}_\ell\|_2$.

For fixed r, candidate generation maximizes $A_{n,t,j,r,\rho}(\boldsymbol{x}) = \widehat{\alpha}_{n,M}(\boldsymbol{x}, r) D_\rho(\boldsymbol{x})$, with

$$\nabla_x A_{n,t,j,r,\rho} = D_\rho \nabla_{\boldsymbol{x}} \widehat{\alpha}_{n,M} + \widehat{\alpha}_{n,M} \nabla_x D_\rho,$$

where $\nabla_{\boldsymbol{x}} \widehat{\alpha}_{n,M}$ is given by Equation (14). Optimization is performed separately for each $r \in \mathcal{R}$. The resulting candidates are reranked by the hard minimum criterion in Equation (13), replacing the sample-average acquisition by $\widehat{\alpha}_n^{\mathrm{quad}}$. The selected design point is assigned $r_n^\star(\boldsymbol{x}) \in \arg\max_{r \in \mathcal{R}} \widehat{\alpha}_n^{\mathrm{quad}}(\boldsymbol{x}, r)$, subject to the serial cost and tie rules.

E Replicate-Aware Mean–Variance BO Algorithm

Algorithm 1 summarizes the serial and parallel policies using the acquisition calculations in Appendix D.

Algorithm 1 Serial or parallel replicate-aware mean–variance BO

Require: Initial design size n_0, initial replication count $r_0 \geq 2$, replication menu $\mathcal{R}$, quantile level τ, batch size $q \geq 1$, reference point $\boldsymbol{r}_{\mathrm{ref}}$, and simulation budget B.

Generate n_0 space-filling design points, collect r_0 replications at each design, and keep all simulator outputs.
Set the number of simulator runs used to $N_{\mathrm{sim}} = n_0 r_0$.
while the total simulation budget has not been reached **do**
    Recompute the mean–variance estimates and structured covariance estimates at all evaluated designs.
    Fit the hierarchical multi-output GP and construct $\widehat{\mathbf{Q}}_{n,r}(\boldsymbol{x})$ for $\boldsymbol{x} \in \mathcal{X}$ and $r \in \mathcal{R}$.
    Construct the current quantile-based Pareto front $\mathcal{P}_n^{(\tau)}$ using Equation (6).
    if $q = 1$ **then**
        Select the design–replication pair $(\boldsymbol{x}_{\mathrm{next}}, r_{\mathrm{next}})$ according to the serial HEQHVI rule in Equation (9).
        Set $\mathcal{A} = \{(\boldsymbol{x}_{\mathrm{next}}, r_{\mathrm{next}})\}$.
    else
        Select the primary pair using the serial HEQHVI rule.
        Select the remaining $q - 1$ pairs sequentially using the HEQHVI–MED criterion in Equation (11) and Equation (13).
        Let $\mathcal{A}$ contain the resulting q pairs.
    end if
    Evaluate every pair $(\boldsymbol{x}, r) \in \mathcal{A}$ and keep all resulting simulator outputs.
    At each reevaluated design point, pool the new outputs with all previously available simulator outputs.
    Increase N_{sim} by the number of newly completed simulator runs.
end while
Refit the surrogate using all available simulator outputs.
return the estimated mean–variance Pareto solution set.

F Experimental Details

F.1 Synthetic Benchmarks

Branin and Hartmann have dimensions $p = 2$ and $p = 6$, respectively, with inputs rescaled to $[0, 1]^p$. We use a smooth, spatially varying noise scale $s_Y(\boldsymbol{x})$ that tends to be larger in regions with lower benchmark response, while allowing local variation beyond this overall relationship. Replicated outputs are generated as

$$Y(\boldsymbol{x}) = G(\boldsymbol{x}) + s_Y(\boldsymbol{x})\varepsilon,$$

so the latent objectives are $(G(\boldsymbol{x}), s_Y^2(\boldsymbol{x}))^\top$. We consider $\varepsilon \sim \mathcal{N}(0, 1)$ and $\varepsilon = (\xi - 4)/2$, where $\xi \sim \mathrm{Gamma}(4, 1)$. Both errors have zero mean and unit variance; thus, for each benchmark, the two output laws share the same latent Pareto front. The two benchmarks and two output laws define four scenarios.

Each scenario is evaluated over 100 independent optimization runs. A run begins with a maximin Latin hypercube of $n_0 = 10p$ design points and five raw replications per design point, followed by 50 BO iterations. Within each run, all policies use the same initial design points and, when their initial replication counts agree, the same simulator outputs. Adaptive serial HEQHVI selects $r \in \{5, 10, 15\}$, while fixed-replication HEQHVI uses $r = 5$. Parallel HEQHVI–MED selects five design points per iteration and assigns five new replications to each design point. All raw simulator outputs are kept and pooled whenever a design is reevaluated.

The comparison methods are GPS-C (Wang et al., 2025), scalar expected improvement for $\mu + v$, serial NEHVI, and parallel qNEHVI. GPS-C models the mean response and uses one simulator output per selected design point. The scalar policy fits a noisy GP to $\bar{Y} + S^2$, which is unbiased for $\mu + v$, and uses five replications per selected design point. Its observation-noise estimate includes the covariance between $\bar{Y}$ and S^2. The NEHVI and qNEHVI policies fit separate noisy GPs to the sample mean and sample variance, using replication-based estimates of their marginal observation variances. NEHVI selects one design point per iteration, whereas qNEHVI selects a batch of five

design points by integrating over the uncertain objective values at both observed and pending designs. Both use five replications per selected design point. GPS-C uses the first output at each initial design point; the remaining policies use all five initial outputs.

For each benchmark, an offline numerical approximation to the reference front $\mathcal{P}_{\mathrm{ref}}$ is constructed from Sobol points. The objective offsets, scales, and reference point $\boldsymbol{r}_{\mathrm{ref}}$ are fixed before the methods are compared. Let $\widehat{\mathcal{P}}_t$ contain the nondominated true objective vectors at distinct designs visited by the end of BO iteration t, expressed on the fixed normalized objective scale. We compute $\mathrm{HV}_t = \mathrm{HV}(\widehat{\mathcal{P}}_t; \boldsymbol{r}_{\mathrm{ref}})$, together with IGD^+ against $\mathcal{P}_{\mathrm{ref}}$. True objectives are used only for evaluation, and both metrics are reported against BO iterations.

Table 1 summarizes the final Pareto-front quality on Hartmann under the two output distributions. Adaptive replication selection consistently improves both hypervolume and IGD^+ over fixed replication, while the serial HEQHVI variants also outperform the other serial methods. The parallel HEQHVI–MED method achieves the best results under both Gaussian and Gamma outputs, indicating an additional benefit from parallel, diversity-aware design selection.

Table 1: Final quality of the Pareto front on Hartmann under Gaussian and Gamma simulation outputs, averaged over 100 independent optimization runs. Higher hypervolume and lower IGD^+ indicate better performance.

	Gaussian		Gamma	
Method	Hypervolume	IGD^+	Hypervolume	IGD^+
HEQHVI adaptive-r	0.6764	0.0120	0.6796	0.0124
HEQHVI fixed-r	0.6619	0.0132	0.6615	0.0206
HEQHVI–MED	**0.7779**	**0.0111**	**0.7774**	**0.0107**
GPS-C	0.3192	0.6330	0.2839	0.7003
SMV	0.4495	0.3733	0.4435	0.3589
NEHVI	0.6369	0.0141	0.6263	0.0261
qNEHVI	0.7527	0.0127	0.7410	0.0137

F.2 F119 Aero-Engine Study

The decision vector is $\boldsymbol{x} = (H, M)^\top$, where altitude $H \in [6, 12]$ and Mach number $M \in [0.5, 1.0]$. The dynamic F119 turbofan model is implemented in MATLAB/Simulink with T-MATS (Chapman et al., 2014a) and returns gross thrust $T(\boldsymbol{x})$. The study uses ten independent optimization runs, each comprising 25 BO iterations. All policies share a fixed 20-point maximin Latin hypercube with five initial outputs per design point. Adaptive HEQHVI selects $r \in \{5, 10, 15\}$, while fixed serial and parallel methods use five replications per selected design point. Parallel methods select five design points per iteration. A fixed empirical reference front is constructed before optimization from estimated mean–variance objective vectors at 128 design locations. All methods use the same fixed reference point for hypervolume and this empirical reference front for IGD^+.

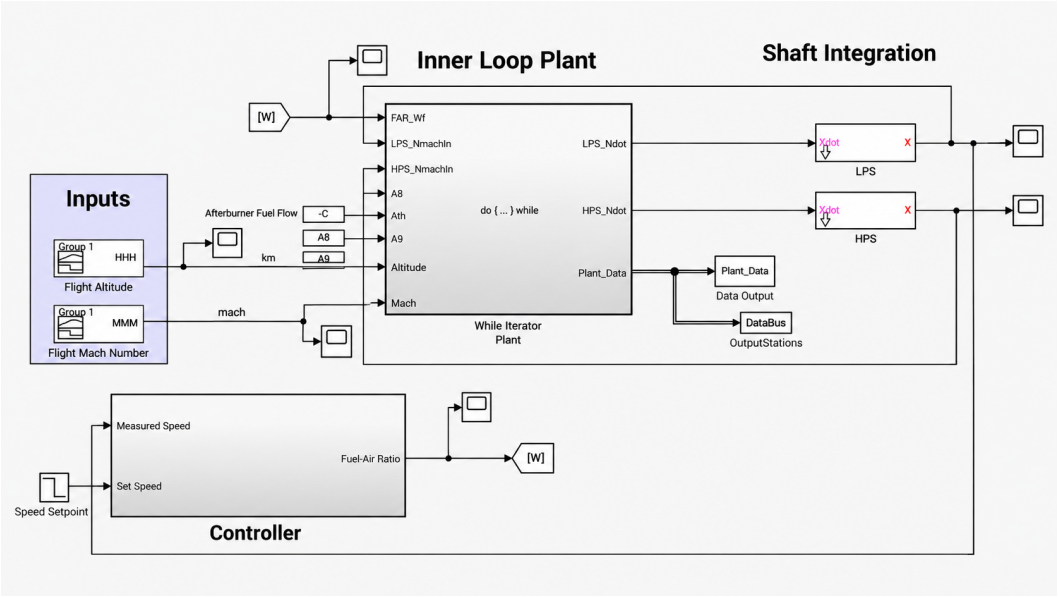


Figure 4: F119 turbofan engine simulation used in the case study.