ADAPTIVE KAPPASHARP: CONDITION-NUMBER SHAPING FOR PREFERENTIAL BAYESIAN OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Preferential Bayesian optimization (PBO) optimizes objectives accessible only through pairwise user comparisons. The standard approach fits a Gaussian process surrogate for observed pairwise comparisons (PairwiseGP) using the Laplace approximation and selects queries with the Expected Utility of Best Option (EUBO) acquisition function. EUBO queries new candidates at each step, producing pairs that share no candidates with previous queries. Each such pair forms an isolated component in the comparison graph, removing one degree of freedom from the likelihood Hessian and making it rank-deficient. This deficiency is structural and cannot be resolved by changing the surrogate modeling approach. Existing approaches to remedy this issue either waste query budget by forcing comparisons to stay connected, or apply uniform regularization that also perturbs directions already well-constrained by the observed comparisons. We propose KappaSharp that enables a diagonal correction to the Hessian to reduce its condition number, with larger corrections where the prior uncertainty is higher. The correction is only applied in the model fitting step, not query selection. An adaptive variant of KappaSharp is also presented that activates the correction only when the surrogate is confident about recent comparisons, avoiding unnecessary corrections when the problem is well-conditioned. On 11 benchmarks (5–20 dimensions), including a 16-dimensional controller tuning problem in plasma medicine, Adaptive KappaSharp outperforms the standard PBO baseline, with up to $+10.9\%$ (p=0.003).

1 INTRODUCTION

Bayesian optimization (BO) (Brochu et al., 2010) is a popular approach for optimizing expensive black-box functions, but requires a scalar objective. Preferential BO (PBO) removes this requirement by learning from pairwise user comparisons (Chu & Ghahramani, 2005; González et al., 2017), with applications in exoskeleton gait optimization (Tucker et al., 2020), materials design (Mikkola et al., 2020), and model-based control design (Shao et al., 2025). Standard PBO approaches fit a Gaussian process surrogate for observed pairwise comparisons (PairwiseGP) via the Laplace approximation (Chu & Ghahramani, 2005) and select queries using the Expected Utility of Best Option (EUBO) acquisition function (Lin et al., 2022; Astudillo et al., 2023), as implemented in BoTorch (Balandat et al., 2020). Recent work has established regret bounds for PBO (Xu et al., 2024), but potential structural limitations of the standard querying strategy have not been investigated.

How queries are selected directly affects the Hessian that PairwiseGP uses to construct its posterior. In the Laplace approximation, the posterior covariance is the inverse of the Hessian of the negative log-posterior (Rasmussen & Williams, 2006). Hence, when the likelihood Hessian is rank-deficient, the full Hessian can become ill-conditioned and, consequently, the posterior can be poorly determined in certain directions. The problem is that observed pairwise comparisons are inherently limited: each comparison only informs the surrogate that one candidate's utility is likely higher than another's. That is, if A,B are compared in one query and C,D in another, the two utility gaps are learned independently. All comparisons can be represented as a *comparison graph*, where nodes are candidates and edges are comparisons. Candidates connected through shared comparisons form *connected components*; candidates in different components have no comparisons connecting them. We show that each disconnected component introduces one zero eigenvalue in the likelihood Hessian, and that these directions are not informed by the observed comparisons. With many disconnected components, the likelihood Hessian becomes severely rank-deficient and the full Hessian can become ill-conditioned.

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.

This becomes a practical challenge because most acquisition functions used in PBO, including EUBO, select two candidates independently per query. When the decision variables are continuous, EUBO maximizes over the full input space, so the selected candidates almost never coincide with previously queried points (Lin et al., 2022). This means each query can add a new disconnected component to the comparison graph, rather than connecting existing ones. Thus, the number of components grows with every query and the rank deficiency of likelihood Hessian worsens. Existing approaches either force the comparison graph to stay connected by constraining query selection (Xu et al., 2024), which limits exploration, or add uniform regularization to Hessian, which addresses the ill-conditioning at the expense of perturbing directions where the observed comparisons provide sufficient information.

Here, we take a different approach: instead of altering the disconnected comparison graph, we look to modify the Hessian of PairwiseGP selectively. We propose KappaSharp, which improves conditioning of the Hessian by shaping its condition number (κ). Rather than regularizing all Hessian entries uniformly, a correction is added to its diagonal entries. The correction is larger where the prior uncertainty is high, and smaller where it is small. The correction directly affects the maximum a posteriori (MAP) estimation of the PairwiseGP posterior; not the query selection step. We additionally present an adaptive variant of KappaSharp that adaptively determines whether the correction must be applied based on PairwiseGP's confidence inferred from recent comparisons. Figure 1 provides an overview of the comparison graph fragmentation problem and how KappaSharp addresses this challenge. KappaSharp does not modify the acquisition function or the querying strategy; thus, it can be applied as a drop-in addition to any PBO approach with minimal computational overhead.

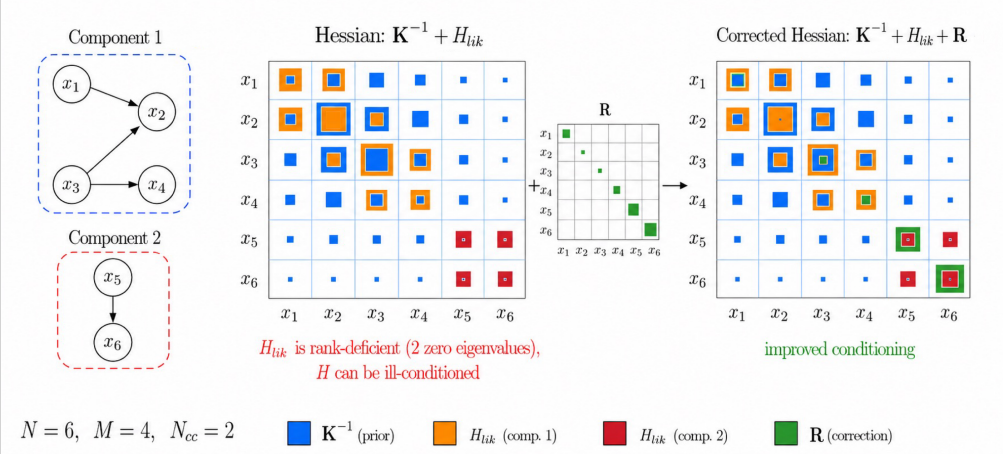

Figure 1: **Left:** An example comparison graph with N=6 candidates, M=4 pairwise comparisons, and N_{cc}=2 disconnected components is shown alongside the corresponding Hessian $\mathbf{H} = \mathbf{K}^{-1} + H_{\text{lik}}$ and corrected Hessian $\mathbf{H} + \mathbf{R}$. Arrows indicate preference direction. **Middle:** The Hessian $\mathbf{H} = \mathbf{K}^{-1} + H_{\text{lik}}$. Blue cells show the prior precision $\mathbf{K}^{-1}$ (dense, all entries). Orange and red cells show the likelihood Hessian H_{lik}, which contributes only within each connected component. The cross-component entries receive no contribution from the observed comparisons, so H_{lik} is rank-deficient and $\mathbf{H}$ can be ill-conditioned. **Right:** The corrected Hessian $\mathbf{H} + \mathbf{R}$. Green diagonal cells show the KappaSharp correction $r_i = \eta^2/(\eta + (\mathbf{K}^{-1})_{ii})$, which is larger where the prior precision is low. Square sizes reflect actual computed values for this example, illustrating that the correction is strongest where $\mathbf{K}^{-1}$ is weakest.

Contributions. The main contributions of this work are as follows:

- C1 We prove that standard PBO querying fragments the comparison graph, making the likelihood Hessian rank-deficient. This is a characteristic of the disconnected comparison graph and cannot be resolved by using alternative surrogate modelling approaches.
- C2 We propose KappaSharp, a diagonal correction to the Hessian applied in MAP estimation of the PairwiseGP posterior, with larger corrections where the prior precision is low.
- C3 We present an adaptive variant of KappaSharp that determines whether to apply the correction based on PairwiseGP's confidence inferred from recent comparisons.
- C4 We evaluate the performance of KappaSharp on 11 benchmarks (5–20 dimensions) with 5 initializations and 60 seeds each, including a dose delivery problem in plasma medicine ,

in which PBO equipped with Adaptive KappaSharp improves over standard PBO without introducing losses.

2 Problem Setup

We aim to find $x^* \in \arg\max_{x \in \mathcal{X}} f(x)$ for a latent utility $f : \mathcal{X} \to \mathbb{R}$, $\mathcal{X} \subseteq \mathbb{R}^d$, that cannot be evaluated directly. At each iteration $t \in \{1, \dots, T\}$, where T is the total number of queries, we select a pair of candidates $(x_{t,a}, x_{t,b}) \in \mathcal{X} \times \mathcal{X}$ and observe only a binary preference. The user does not observe the utility values; instead, the preference is generated through noisy comparisons:

$$y_{t,a} = f(x_{t,a}) + \epsilon_{t,a}, \quad y_{t,b} = f(x_{t,b}) + \epsilon_{t,b}, \quad \epsilon_{t,a}, \epsilon_{t,b} \overset{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2), \tag{1}$$

with noise scale $\sigma > 0$. The user reports $x_{t,a} \succ x_{t,b}$ if $y_{t,a} > y_{t,b}$. This yields the probit likelihood

$$P(x_{t,a} \succ x_{t,b} \mid \mathbf{f}) = \Phi\left(\frac{f(x_{t,a}) - f(x_{t,b})}{\sigma\sqrt{2}}\right), \tag{2}$$

where $\mathbf{f}$ denotes the vector of latent utilities at all queried candidates and Φ is the standard normal cumulative distribution function (CDF). Each observation is recorded as an ordered pair (x_t^+, x_t^-), where x_t^+ is the preferred candidate.

3 Related Work

Preferential Bayesian Optimization. Chu & Ghahramani (2005) introduced GP-based preference learning, and González et al. (2017) formulated PBO as a sequential optimization approach. The EUBO acquisition function (Lin et al., 2022) and its batch variant qEUBO (Astudillo et al., 2023) are currently the standard choice for query selection under the Laplace approximation. Xu et al. (2024) provided the first regret bounds for PBO by enforcing a connected comparison graph, at the cost of reduced exploration. Alternative posterior approximations have also been proposed. Benavoli et al. (2020) derived the exact skew-GP posterior; Takeno et al. (2023) developed a practical implementation of SkewGP; and Mikkola et al. (2020) proposed Gibbs sampling as an alternative. None of these developments addresses the rank deficiency problem identified in this work: when the comparison graph is disconnected, the observed comparisons provide no information linking different components (Proposition 2), regardless of the surrogate modelling approach. SkewGP also requires substantially more computation per posterior evaluation than the Laplace approximation.

Comparison Graphs and Spectral Structure. The connection between pairwise comparisons and graph Laplacians is investigated extensively in spectral ranking (Jiang et al., 2011). Shah et al. (2016) showed that the Fisher information matrix for pairwise comparison models inherits the Laplacian structure. The fact that the probit Hessian takes a Laplacian form has been known since the work of Chu & Ghahramani (2005), but its impact on the PairwiseGP posterior under standard querying has not been analyzed. We show that EUBO with q=2 fragments the comparison graph into many disconnected components, causing the likelihood Hessian H_{lik} to lose rank, which will in turn make the full Hessian ill-conditioned (Section 4).

Posterior Stabilization in Gaussian Processes. Hartmann & Vanhatalo (2019) replaced the Hessian with the Fisher information matrix for GP classification. Kündig & Sigrist (2024) introduced preconditioners for Vecchia-Laplace approximations. In continual learning, Elastic Weight Consolidation (EWC) (Kirkpatrick et al., 2017) and online Laplace (Immer et al., 2021) add diagonal regularization based on the curvature of the loss. These methods do not account for the graph-Laplacian rank structure that arises in pairwise comparison models. When H_{lik} is rank-deficient, the Hessian does not receive contribution from the observed comparisons in certain directions. Thus, the eigenvalues of the Hessian are determined solely by the prior precision, which can be much smaller than in directions where H_{lik} contributes. The proposed correction scheme addresses this imbalance by shaping the condition number of the Hessian.

4 STRUCTURAL ANALYSIS

The iterative PBO procedure starts from N_0 initial candidate points $x_1, \dots, x_{N_0} \in \mathcal{X}$ with M_0 comparisons among them, forming $N_{\mathrm{cc},0}$ connected components in the comparison graph. As discussed in Section 1, under EUBO with batch size $q{=}2$ (two candidate points per query), each subsequent query selects two new candidate points independent of all previous ones. After $t{-}1$ such queries, we have $N_{t-1} = N_0 + 2(t{-}1)$ candidates and $M_{t-1} = M_0 + (t{-}1)$ comparisons. With a slight abuse of notation, we write $N = N_{t-1}$ and $M = M_{t-1}$ for the number of candidates and comparisons at iteration t, respectively.

A GP prior $\mathbf{f} \sim \mathcal{N}(\mathbf{0}, \mathbf{K})$ is placed over the latent utility values $\mathbf{f} = (f(x_1), \dots, f(x_N))^\top$, where $\mathbf{K} \in \mathbb{R}^{N\times N}$ with $K_{ij} = k(x_i, x_j)$ for a kernel function k. Accordingly, the m-th comparison records that x_m^+ was preferred over x_m^-. We encode all M comparisons in an incidence matrix $\mathbf{C} \in \mathbb{R}^{M\times N}$, where row m has $+1$ in the column corresponding to x_m^+, -1 in the column corresponding to x_m^-, and zero elsewhere. Let $\mathcal{D} = \{(x_m^+, x_m^-)\}_{m=1}^M$ denote the set of observed pairwise comparisons, where each pair records the preferred and non-preferred candidate. The Laplace approximation gives the posterior $p(\mathbf{f} \mid \mathcal{D}) \approx \mathcal{N}(\hat{\mathbf{f}}, \mathbf{H}^{-1})$, where $\hat{\mathbf{f}} = \arg\min_{\mathbf{f}} -\log p(\mathcal{D} \mid \mathbf{f}) + \frac{1}{2}\mathbf{f}^\top \mathbf{K}^{-1}\mathbf{f}$ is the MAP estimate of $\mathbf{f}$ and

$$\mathbf{H} = \mathbf{K}^{-1} + \underbrace{\mathbf{C}^\top \mathbf{W} \mathbf{C}}_{H_{\mathrm{lik}}} \tag{3}$$

is the posterior precision matrix, with $\mathbf{W} = \mathrm{diag}(\omega_1, \dots, \omega_M)$ a diagonal matrix of positive weights that depend on $\hat{\mathbf{f}}$, as defined in Proposition 1. The likelihood Hessian H_{lik} has a specific structure.

Proposition 1 (H_{lik} is a weighted graph Laplacian)**.** *For every $\mathbf{f} \in \mathbb{R}^N$ with finite comparison scores, the data-dependent component of the Laplace precision, $H_{\mathrm{lik}}(\mathbf{f}) = -\nabla^2_{\mathbf{f}} \log p(\mathcal{D} \mid \mathbf{f})$, satisfies*

$$H_{\mathrm{lik}}(\mathbf{f}) = \mathbf{C}^\top \mathbf{W}(\mathbf{f})\, \mathbf{C}, \tag{4}$$

where $\mathbf{W} = \mathrm{diag}(\omega_1, \dots, \omega_M)$ with $\omega_m = (2\sigma^2)^{-1}\psi(z_m)\big(\psi(z_m) + z_m\big) > 0$, $z_m = \mathbf{C}_m^\top \mathbf{f}/(\sigma\sqrt{2})$, ϕ is the standard normal probability density function (PDF), Φ is the CDF, and $\psi(z) = \phi(z)/\Phi(z)$ is the inverse Mills ratio. The rank of H_{lik} is

$$\mathrm{rank}(H_{\mathrm{lik}}) = N - N_{\mathrm{cc}}, \tag{5}$$

where N_{cc} is the number of connected components of the comparison graph.

Within each connected component, H_{lik} constrains relative utility differences. That is, the likelihood cannot determine how utilities in one component relate to utilities in another. Proofs for all propositions and theorems in this section and Section 5 are given in Section A.

Proposition 2 (Likelihood-null directions)**.** *Let $\ker(\mathbf{C}) = \{\mathbf{f} \in \mathbb{R}^N : \mathbf{C}\mathbf{f} = \mathbf{0}\}$ be the null space of $\mathbf{C}$ and $\mathrm{range}(\mathbf{C}^\top)$ be its orthogonal complement. Let $\mathbf{Q}_0 \in \mathbb{R}^{N\times N_{\mathrm{cc}}}$ and $\mathbf{Q}_1 \in \mathbb{R}^{N\times(N-N_{\mathrm{cc}})}$ be orthonormal bases for $\ker(\mathbf{C})$ and $\mathrm{range}(\mathbf{C}^\top)$, respectively. Decompose the latent utility vector as $\mathbf{f} = \mathbf{Q}_0\xi + \mathbf{Q}_1\zeta$, where $\xi \in \mathbb{R}^{N_{\mathrm{cc}}}$ are the null-space coordinates and $\zeta \in \mathbb{R}^{N-N_{\mathrm{cc}}}$ are the range-space coordinates. Then:*

- *(a) The likelihood does not depend on ξ: $p(\mathcal{D} \mid \xi, \zeta) = p(\mathcal{D} \mid \zeta)$.*
- *(b) Under the Laplace approximation, the posterior over ξ given ζ equals the prior: $p_L(\xi \mid \zeta, \mathcal{D}) = p_0(\xi \mid \zeta)$.*

The null-space coordinates ξ represent a constant shift in the utility of all candidates within the same connected component, i.e., a group of candidate points x that are linked through a chain of pairwise comparisons. Since comparisons only measure utility *differences* within a pair, such shifts do not change any comparison outcome, and the posterior over ξ is determined entirely by the GP prior.

Comparison graph fragmentation under EUBO. Given the posterior $\mathcal{N}(\hat{\mathbf{f}}, \mathbf{H}^{-1})$, the EUBO acquisition function with $q{=}2$ selects the pair (x_a, x_b) that maximizes the expected utility of the best option in the pair. For a pair with posterior mean difference $\delta = \hat{f}(x_a) - \hat{f}(x_b)$ and predictive standard deviation $s = \sqrt{\mathrm{Var}[f(x_a) - f(x_b) \mid \mathcal{D}]}$, the EUBO score is then

$$\mathrm{EUBO}(x_a, x_b) = \delta\, \Phi(\delta/s) + s\, \phi(\delta/s), \tag{6}$$

In continuous decision domains, these candidates are almost always new points (Lin et al., 2022). Thus, after T queries from N_0 initial points with $N_{\mathrm{cc},0}$ connected components, there are $N_0 + 2T$ candidates in up to $N_{\mathrm{cc},0} + T$ components.

We now quantify how much the PairwiseGP surrogate can learn from the observed pairwise comparisons using the *effective degrees of freedom* $d_{\mathrm{eff}} \in [0, N)$. Let $\tilde{\mathbf{M}} = \mathbf{K}^{1/2} H_{\mathrm{lik}} \mathbf{K}^{1/2}$, which captures how much information the observed comparisons add beyond the prior in each direction of the latent utility space. Each eigenvalue μ_i of $\tilde{\mathbf{M}}$ corresponds to one such direction, and $\mu_i/(1+\mu_i)$ is close to 1 when the comparisons dominate and close to 0 when the prior dominates. $d_{\mathrm{eff}} = \sum_i \mu_i/(1+\mu_i)$ counts how many directions are effectively informed by the observed comparisons.

Proposition 3 (Effective information dimension)**.** *Given* $\tilde{\mathbf{M}}$ *and* d_{eff} *as defined above:*

(a) $d_{\mathrm{eff}} \le \mathrm{rank}(H_{\mathrm{lik}}) = N - N_{\mathrm{cc}}$*, with equality only in the limit* $\mu_i \to \infty$*.*

(b) Under EUBO $q{=}2$*, assume each iteration queries a comparison between two previously unseen points and the initial* N_0 *points form* $N_{\mathrm{cc},0}$ *connected components. Then,* $N = 2T + N_0$*,* $N_{\mathrm{cc}} = N_{\mathrm{cc},0} + T$*, and* $d_{\mathrm{eff}}/N \le (N_0 + T - N_{\mathrm{cc},0})/(N_0 + 2T)$*, which approaches* $1/2$ *as* $T \to \infty$*.*

(c) Under connected querying ($N_{\mathrm{cc}}{=}1$*):* $d_{\mathrm{eff}} \le N - 1$*.*

(d) With any positive-definite augmentation $\mathbf{R} \succ 0$ *satisfying* $\mathbf{K}^{1/2}\mathbf{R}\mathbf{K}^{1/2} \succeq \beta\mathbf{I}$ *for some scalar* $\beta > 0$*, define* $d_{\mathrm{eff}}^{\mathrm{aug}}$ *as the effective degrees of freedom computed from* $\mathbf{H} + \mathbf{R}$ *(i.e., replacing* $\tilde{\mathbf{M}}$ *by* $\mathbf{K}^{1/2}(H_{\mathrm{lik}} + \mathbf{R})\mathbf{K}^{1/2}$*). Then:*

$$d_{\mathrm{eff}}^{\mathrm{aug}} \ge d_{\mathrm{eff}} + \frac{N_{\mathrm{cc}}\beta}{1+\beta}. \tag{7}$$

In Proposition 3, part (a) indicates that each disconnected component removes one effective degree of freedom. Part (b) indicates that under standard EUBO, at most half the degrees of freedom are informed by the observed pairwise comparisons. Part (c) indicates that connected querying removes this ceiling. Part (d) indicates that adding a positive-definite matrix to the Hessian can increase d_{eff} algebraically, but this does not represent real information from comparisons (see Section 5).

5 Condition-Number Shaping of Hessian

In standard PBO, the posterior of PairwiseGP is estimated via MAP estimation in every iteration. Let $J_0(\mathbf{f}) = -\log p(\mathcal{D} \mid \mathbf{f}) + \frac{1}{2}\mathbf{f}^\top \mathbf{K}^{-1}\mathbf{f}$ denote the negative log-posterior up to a constant, and let $\hat{\mathbf{f}}_0 = \arg\min_{\mathbf{f}} J_0(\mathbf{f})$ be the standard MAP estimate. As shown in Section 4, the likelihood Hessian H_{lik} is rank-deficient when the comparison graph is disconnected, which can make the full Hessian $\mathbf{H} = \mathbf{K}^{-1} + H_{\mathrm{lik}}(\hat{\mathbf{f}}_0)$ ill-conditioned. We propose to replace the standard MAP estimation with the augmented problem

$$\hat{\mathbf{f}}_\eta = \arg\min_{\mathbf{f}} \; J_0(\mathbf{f}) + \tfrac{1}{2}\mathbf{f}^\top \mathbf{R}(\eta)\,\mathbf{f}, \tag{8}$$

where $\mathbf{R}(\eta) = \mathrm{diag}(r_1(\eta), \ldots, r_N(\eta))$ is a diagonal correction matrix that adds a penalty to each diagonal entry of the Hessian. The scalar $\eta > 0$ controls the overall correction strength. After solving for $\hat{\mathbf{f}}_\eta$, we evaluate the Hessian of J_0 (not the augmented objective) at $\hat{\mathbf{f}}_\eta$, denoted as $\mathbf{H}_\eta = \mathbf{K}^{-1} + H_{\mathrm{lik}}(\hat{\mathbf{f}}_\eta)$, and use its inverse as the Laplace covariance: $\Sigma_{\mathrm{MO}} = \mathbf{H}_\eta^{-1}$. That is, $\mathbf{R}$ shifts the MAP from $\hat{\mathbf{f}}_0$ to $\hat{\mathbf{f}}_\eta$ but does not enter the covariance. We describe the design of $\mathbf{R}$ next and justify the deployment of MAP-only in the following.

5.1 KappaSharp Correction

Each diagonal entry $r_i(\eta)$ depends on the prior precision $(\mathbf{K}^{-1})_{ii}$ at the i-th candidate point.

Proposition 4 (Correction properties)**.** *Define* $a_i = (\mathbf{K}^{-1})_{ii} > 0$*, which denotes the prior precision at the* i*-th candidate point, and let*

$$r_i(\eta) = \frac{\eta^2}{\eta + a_i}. \tag{9}$$

Then, for $\eta > 0$:

(a) $\partial r_i / \partial a_i < 0$: larger prior precision → less correction;

(b) $0 < r_i < \eta$: the correction is bounded;

(c) as $a_i \to 0$: $r_i \to \eta$ (maximum correction where the prior is weak);

(d) as $a_i \to \infty$: $r_i \to 0$ (no correction where the prior already dominates).

The correction uses only the diagonal entries $(\mathbf{K}^{-1})_{ii}$, not the likelihood Hessian H_{lik} directly. Adding $\mathbf{R}$ to the Hessian is equivalent to strengthening the prior precision from $\mathbf{K}^{-1}$ to $\mathbf{K}^{-1} + \mathbf{R}$. The ill-conditioning of $\mathbf{H} = \mathbf{K}^{-1} + H_{\text{lik}}$ is worst in directions where H_{lik} contributes nothing and $\mathbf{K}^{-1}$ is weak. Under EUBO querying in continuous domains, these two conditions tend to occur simultaneously: newly queried points that form isolated components in the comparison graph tend to have low kernel correlation with existing points, so both their likelihood contribution and their prior uncertainty are high. By applying larger corrections where $(\mathbf{K}^{-1})_{ii}$ is small, $\mathbf{R}$ boosts the weakest eigenvalues of $\mathbf{H}$, improving its conditioning without requiring direct access to the structure of H_{lik}. The overall correction strength η^* is calibrated using a condition-number target on the full Hessian $\mathbf{H} + \mathbf{R}$, which accounts for H_{lik} indirectly. The condition number $\kappa(\mathbf{H}) = \lambda_{\max}(\mathbf{H})/\lambda_{\min}(\mathbf{H})$ measures the ratio between the largest and smallest eigenvalues of $\mathbf{H}$. A large $\kappa(\mathbf{H})$ indicates that some directions are much better determined than others. We choose η^* as the smallest η that reduces $\kappa(\mathbf{H} + \mathbf{R}(\eta))$ to a target level, which we call κ-calibration.

Proposition 5 (κ-calibration is well-posed)**.** *Let $\mathbf{H} = \mathbf{K}^{-1} + H_{\text{lik}}(\hat{\mathbf{f}}_0)$ be the Hessian at $\hat{\mathbf{f}}_0$ and $\alpha \in (0, 1)$. The correction $\mathbf{R}(\eta)$ is continuous, positive semi-definite, and monotonically increasing in η (i.e., $\eta' > \eta \Rightarrow \mathbf{R}(\eta') \succeq \mathbf{R}(\eta)$); $\kappa(\mathbf{H} + \mathbf{R}(\eta)) \to 1$ as $\eta \to \infty$; and therefore*

$$\eta^* = \min_{\eta} \left\{ \eta \geq 0 : \kappa(\mathbf{H} + \mathbf{R}(\eta)) \leq \kappa(\mathbf{H})^{1-\alpha} \right\} \tag{10}$$

exists and is finite.

We set $\alpha = 0.1$ throughout, targeting $\kappa(\mathbf{H})^{0.9}$. A sensitivity analysis over α is shown in Section I.

Bounded MAP shift. If the correction shifts the MAP too far from the original estimate, the resulting posterior may no longer reflect the observed pairwise comparisons. The following proposition bounds the shift $\Delta = \hat{\mathbf{f}}_\eta - \hat{\mathbf{f}}_0$.

Proposition 6 (Bounded MAP shift)**.** *The shift $\Delta = \hat{\mathbf{f}}_\eta - \hat{\mathbf{f}}_0$ satisfies $\|\Delta\|_{\bar{\mathbf{H}}} \leq \frac{\mu_{\max}}{1+\mu_{\max}} \left\|\hat{\mathbf{f}}_0\right\|_{\bar{\mathbf{H}}}$, where $\|\cdot\|_{\bar{\mathbf{H}}}$ is the norm induced by the path-averaged Hessian $\bar{\mathbf{H}}$ of J_0, and $\mu_{\max}$ is the largest eigenvalue of $\bar{\mathbf{H}}^{-1/2}\mathbf{R}\bar{\mathbf{H}}^{-1/2}$ (the path-averaged Hessian and its norm are defined in the proof in Section A).*

Note that shifting the MAP changes the weights $\mathbf{W}$ but not the rank of H_{lik}: since $\mathbf{W}(\mathbf{f}) \succ 0$ for all finite $\mathbf{f}$, $\text{rank}(H_{\text{lik}}(\mathbf{f})) = N - N_{\text{cc}}$ regardless of $\mathbf{f}$ (see Proposition 10 in Section A)).

Why $\mathbf{R}$ should not enter the covariance. Recall from Proposition 2 that the observed pairwise comparisons provide no information in the null-space directions $\ker(\mathbf{C})$; thus, the posterior covariance in those directions should equal the prior. If we were to include $\mathbf{R}$ in the covariance, the added $\mathbf{R}$ would shrink the posterior covariance in the null-space directions to values smaller than the prior covariance. PairwiseGP would then underestimate its uncertainty in those directions, and EUBO would stop exploring them. By excluding $\mathbf{R}$ from the covariance, MAP-only preserves the correct uncertainty in the null space.

Theorem 7 (Pseudo-confidence in null directions)**.** *Using the decomposition $\mathbf{f} = \mathbf{Q}_0\xi + \mathbf{Q}_1\zeta$ from Proposition 2, let $\mathbf{A} = \mathbf{Q}_0^\top \mathbf{K}^{-1}\mathbf{Q}_0$ be the prior precision on the null-space coordinates ξ. Full augmentation $(\mathbf{H}_\eta + \mathbf{R})^{-1}$ shrinks the conditional covariance of ξ, given ζ, to $(\mathbf{A} + \mathbf{Q}_0^\top \mathbf{R}\mathbf{Q}_0)^{-1} \prec \mathbf{A}^{-1}$ since $\mathbf{Q}_0^\top \mathbf{R}\mathbf{Q}_0 \succ 0$. MAP-only preserves $\mathbf{A}^{-1}$, matching the exact posterior (Proposition 2).*

Full augmentation also reduces EUBO scores for all candidate pairs, which discourages exploration.

Proposition 8 (Full augmentation contracts EUBO scores)**.** *Let $\Sigma_{\text{aug}} := (\mathbf{H}_\eta + \mathbf{R})^{-1}$. For any candidate pair (x_a, x_b), let $s^2_\Sigma(a, b)$ denote the predictive variance of the utility difference $f(x_a) -$*

$f(x_b)$ under covariance Σ. Since $\Sigma_{\text{aug}} \preceq \Sigma_{\text{MO}}$, we have $s^2_{\text{aug}}(a,b) \leq s^2_{\text{MO}}(a,b)$ for every pair. Since q=2 EUBO is strictly increasing in the predictive standard deviation at fixed mean, the EUBO score under full augmentation is lower than under MAP-only for every candidate pair:

$$\text{EUBO}_{\text{aug}}(a,b) \leq \text{EUBO}_{\text{MO}}(a,b) \quad \textit{for every pair } (x_a, x_b). \tag{11}$$

The cost of MAP-only is bounded. Excluding $\mathbf{R}$ from the covariance avoids both issues. The acquisition gap from between the MAP-only and full-augmentation strategies is bounded.

Theorem 9 (Bounded acquisition gap)**.** *The difference in EUBO between the MAP-only maximizer and the full-augmentation maximizer is at most $O(\lambda_1)$, where λ_1 is the largest eigenvalue of $\mathbf{H}_\eta^{-1/2}\mathbf{R}\,\mathbf{H}_\eta^{-1/2}$, measuring how large the correction $\mathbf{R}$ is relative to the Hessian $\mathbf{H}_\eta$. The precise bound is given in Section A and is finite for any continuous kernel on a compact domain.*

5.2 Adaptive Activation

Applying KappaSharp at every iteration, referred to as *Static KS* in our experiments, is beneficial when the rank deficiency of the likelihood Hessian is the primary bottleneck limiting the surrogate. However, it can be detrimental when the prior already provides sufficient conditioning. We observe that PairwiseGP's confidence on recently queried pairs separates these two regimes: when the rank deficiency of the likelihood Hessian dominates, the surrogate quickly becomes confident about individual comparisons; when the prior is already well-conditioned, the surrogate remains uncertain. The adaptive variant of KappaSharp uses this separation to decide when to activate the correction.

Decisiveness score. After fitting the surrogate on $\mathcal{D}_{t-1}$, we measure how confidently it distinguishes a recently compared pair. Let $\mu_{\text{diff}} = |\hat{f}(x^+_{t-1}) - \hat{f}(x^-_{t-1})|$ be the absolute posterior mean difference and σ^2_{diff} the posterior variance of $f(x^+_{t-1}) - f(x^-_{t-1})$. The surrogate's predicted probability that the comparison outcome matches its ranking is $p_t = \Phi\Big(\mu_{\text{diff}}/\sqrt{\sigma^2_{\text{diff}} + 2\sigma^2}\Big) \in [1/2, 1)$, where $2\sigma^2$ is the comparison noise variance from the probit model (Equation (1)). A high value of p_t indicates that the surrogate predicts a clear utility difference between the two candidates, whereas $p_t \approx 0.5$ indicates that the comparison is ambiguous. We define the *decisiveness score* as

$$s_{\text{exp},t} = 1 - \frac{\mathcal{H}(p_t)}{\log 2}, \tag{12}$$

where $\mathcal{H}(p) = -p\log p - (1-p)\log(1-p)$ is the binary entropy. The score ranges from 0 (coin flip) to 1 (decisive).

Activation rule. We track an exponential moving average to smooth the step-to-step variance of $s_{\text{exp},t}$:

$$\bar{s}_t = (1 - \alpha_{\text{ema}})\,\bar{s}_{t-1} + \alpha_{\text{ema}} s_{\text{exp},t}, \quad \alpha_{\text{ema}} = 0.2, \tag{13}$$

initialized at $\bar{s}_0 = 0$. The correction activates at step t when $\bar{s}_t \geq \tau$, with threshold $\tau = 0.30$. In each step, $\tau = 0.30$ corresponds to $p_t \geq 0.81$ (Proposition 13 in Section D). The activation rule uses the EMA $\bar{s}_t$, not the per-step score. Thus, $\bar{s}_t \geq \tau$ indicates sufficient average decisiveness over recent comparisons. Algorithm 1 summarizes the Adaptive KappaSharp approach. When the correction is active, the additional computation consists of reading the diagonal of $\mathbf{K}^{-1}$ to form $\mathbf{R}$, searching over 60 log-spaced values of η to solve the κ-calibration problem, and running one extra Newton solve for the augmented MAP. The better-conditioned Hessian often reduces the number of Newton iterations needed, partially or fully offsetting this cost (Section L). When the correction is inactive, the only overhead per step is evaluating s_{exp} from a single posterior prediction.

6 Experiments

We benchmark Adaptive KappaSharp against the standard PairwiseGP baseline and several alternatives on 11 problems spanning 5 to 20 dimensions. The baselines include Static KS (KappaSharp applied at every iteration without the activation rule), Fixed20 (correction activated only after step 20, without the decisiveness criterion), Connected Querying (one candidate per query is fixed to the previous query point $x'_t = x_{t-1}$, following the querying strategy of Xu et al. (2024); EUBO is retained as the acquisition function for fair comparison), and the unmodified PairwiseGP-Laplace + EUBO baseline (B). We report ablations and diagnostic comparisons with non-Laplace inference.

Algorithm 1 Adaptive KappaSharp for Preferential BO

Require: Initial data $\mathcal{D}_0$, GP prior $\mathbf{K}$, threshold $\tau = 0.30$, calibration $\alpha = 0.1$, EMA rate $\alpha_{\text{ema}} = 0.2$, minimum step $t_{\min} = 8$

Initialize $\bar{s}_0 \leftarrow 0$
for $t = 1, 2, \ldots, T$ **do**
  Fit PairwiseGP on $\mathcal{D}_{t-1}$: compute MAP $\hat{\mathbf{f}}_0$, Hessian $\mathbf{H} = \mathbf{K}^{-1} + H_{\text{lik}}$
  Compute $s_{\text{exp},t}$ from PairwiseGP's prediction on most recent comparison {Eq. (12)}
  Update $\bar{s}_t \leftarrow (1 - \alpha_{\text{ema}})\, \bar{s}_{t-1} + \alpha_{\text{ema}} s_{\text{exp},t}$ {Eq. (13)}
  if $t \geq t_{\min}$ **and** $\bar{s}_t \geq \tau$ **then**
    Compute $a_i = (\mathbf{K}^{-1})_{ii}$ and $\mathbf{R}(\eta) = \text{diag}(\eta^2/(\eta + a_i))$ {Eq. (9)}
    Solve $\eta^* = \min_{\eta} \{\eta \geq 0 : \kappa(\mathbf{H} + \mathbf{R}(\eta)) \leq \kappa(\mathbf{H})^{1-\alpha}\}$ {Eq. (10)}
    Re-solve MAP: $\hat{\mathbf{f}}_\eta \leftarrow \arg\min J_0(\mathbf{f}) + \frac{1}{2}\mathbf{f}^\top \mathbf{R}(\eta^*)\mathbf{f}$ {Eq. (8)}
    Set approximate posterior: $\mathcal{N}(\hat{\mathbf{f}}_\eta,\ \mathbf{H}^{-1}|_{\hat{\mathbf{f}}_\eta})$ {Unaugmented covariance at shifted MAP}
  else
    Use standard posterior: $\mathcal{N}(\hat{\mathbf{f}}_0,\ \mathbf{H}^{-1}|_{\hat{\mathbf{f}}_0})$
  end if
  Select next query (x_a, x_b) via EUBO on the current posterior
  Observe comparison outcome; update $\mathcal{D}_t \leftarrow \mathcal{D}_{t-1} \cup \{(x_t^+, x_t^-)\}$
end for

6.1 Setup

Benchmarks. The 11 benchmarks span 5 to 20 dimensions: DTLZ2-8D, DTLZ2-16D, DTLZ2-20D (multi-objective PBO test functions), Plasma-16D (model-based control design in plasma medicine (Shao et al., 2025)), Ackley-8D, Hartmann-6D, Levy-10D, Levy-20D, Vehicle-5D (single-objective), CarCab-7D (vehicle cabin design), and RobotPush-14D (Box2D physics simulation).

Protocol. All methods use PairwiseGP with the Laplace approximation and EUBO acquisition (Lin et al., 2022), implemented in BoTorch (Balandat et al., 2020). Each run consists of $T = 50$ sequential pairwise queries, with comparison noise $\sigma = 0.1$. We test five initialization conditions: `pool_n6_k15` (6 candidates, all 15 pairwise comparisons; primary), `pool_n8_k15` (8 candidates, 15 comparisons), and three pure-matching conditions `pure_k3`, `pure_k5`, `pure_k15` (k non-overlapping pairs among $2k$ candidates). Each condition uses 60 seeds with paired same-node execution and common random numbers (CRN). The total number of runs in the main grid is $16{,}500$ (11 benchmarks $\times$ 5 inits $\times$ 5 methods $\times$ 60 seeds).

Evaluation. We report percentage improvement in final best utility over baseline B. Significance is assessed via paired t-test ($p < 0.05$, uncorrected). A benchmark-method pair is classified as win (W), loss (L), or neutral (N) based on the sign and significance of the improvement.

6.2 Baselines and ablations

Table 1 compares all methods on the primary initialization (pool_n6_k15, 60 seeds). Static KS wins on 5 benchmarks (DTLZ2-16D $+24.5\%$, DTLZ2-8D $+23.2\%$, DTLZ2-20D $+17.8\%$, Plasma $+9.5\%$, Levy-10D $+18.7\%$) but loses on Ackley-8D (-15.6%) and Hartmann-6D (-4.8%). Connected Querying achieves 0W/7L/4N, with -93.9% on DTLZ2-8D and -61.9% on Ackley-8D (both $p < 0.01$); see Section J for details.

We ablate the correction design in Table 2 (Section B). Uniform jitter ($\mathbf{R} = cI$): 0W/5L. Uniform diagonal ($\mathbf{R} = \eta^* I$, without $(\mathbf{K}^{-1})_{ii}$ weighting): 3W/2L. MetricOnly ($\mathbf{R}$ in the Newton step but J_0 unchanged, so MAP is the same): 1W/0L. CovOnly (MAP unchanged, covariance uses $(\mathbf{H} + \mathbf{R})^{-1}$): 1W/0L. These results indicate that the MAP correction is the main source of improvement. We also compare against EP-EI and Gibbs-EI (Takeno et al., 2023). As discussed in Section 3, changing the inference method does not resolve the rank deficiency of the likelihood Hessian (Proposition 2). Both score worse than Laplace+EUBO (Section K); EP and Gibbs do not have an acquisition function like EUBO; hence, they use EI, designed for scalar observations rather than pairwise comparisons.

Table 1: Benchmark results (pool_n6_k15 init, 60 seeds). B: PairwiseGP-Laplace + EUBO baseline. Static KS: KappaSharp without activation rule. Fixed20: correction on at step $t \geq 20$. Conn. Query: connected querying strategy of Xu et al. (2024), with EUBO retained for fair comparison. Adaptive KS: proposed method with confidence-based activation (τ=0.30). Gain (%) relative to B. **Bold**: best mean. *: $p<0.05$, **: $p<0.01$ (paired t-test).

Benchmark	d	B	Static KS		Fixed20		Conn. Query		Adaptive KS	
		mean±std	mean±std	Gain	mean±std	Gain	mean±std	Gain	mean±std	Gain
DTLZ2-16D	16	-1.031±0.341	**-0.778±0.281**	+24.5%**	-0.783±0.282	+24.0%**	-1.016±0.327	+1.5%	-0.875±0.265	+15.1%**
DTLZ2-8D	8	-0.184±0.134	**-0.142±0.068**	+23.2%*	-0.167±0.114	+9.6%	-0.357±0.189	-93.9%**	-0.163±0.101	+11.7%
DTLZ2-20D	20	-1.439±0.390	**-1.183±0.379**	+17.8%**	-1.256±0.453	+12.7%**	-1.350±0.342	+6.2%	-1.268±0.444	+11.9%*
Plasma-16D	16	-92.785±24.954	-84.005±20.141	+9.5%**	-85.676±22.113	+7.7%*	-103.451±38.509	-11.5%*	**-82.642±21.800**	+10.9%**
Levy-10D	10	-17.193±11.645	-13.985±9.807	+18.7%*	-13.675±8.420	+20.5%**	-20.920±11.493	-21.7%*	**-13.069±9.307**	+24.0%**
Levy-20D	20	-64.416±22.656	-62.420±29.474	+3.1%	-58.497±23.673	+9.2%*	-70.013±22.260	-8.7%	**-58.167±27.639**	+9.7%*
Ackley-8D	8	**-1.168±0.756**	-1.350±0.773	-15.6%*	-1.265±0.786	-8.3%	-1.891±0.776	-61.9%**	-1.184±0.739	-1.4%
Hartmann-6D	6	**2.959±0.341**	2.818±0.444	-4.8%**	2.810±0.513	-5.0%**	2.611±0.586	-11.8%**	2.934±0.456	-0.9%
CarCab-7D	7	-3.888±0.081	-3.900±0.084	-0.3%	-3.894±0.091	-0.2%	-3.895±0.083	-0.2%	**-3.889±0.088**	-0.0%
Vehicle-5D	5	-5038.884±0.261	-5038.994±0.945	-0.0%	-5039.100±0.993	-0.0%	-5039.640±1.875	-0.0%*	**-5038.873±0.136**	+0.0%
RobotPush-14D	14	3.576±1.799	3.639±1.851	+1.8%	3.736±1.665	+4.5%	2.691±1.820	-24.8%**	**3.776±1.810**	+5.6%
W/L/N			5/2/4		5/1/5		0/7/4		5/0/6	

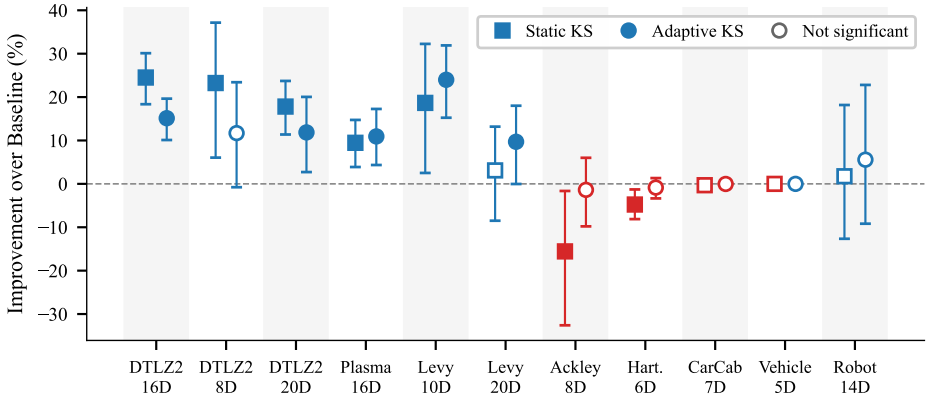

Figure 2: Effect sizes (% improvement over baseline, 95% CI, pool_n6_k15, 60 seeds). Filled markers indicate $p < 0.05$. Adaptive KS: 5W/0L/6N. Static KS: 5W/2L/4N.

6.3 Main results

Adaptive KS scores **5W/0L/6N** on the primary initialization: DTLZ2-16D ($+15.1\%$), DTLZ2-20D ($+11.9\%$), Plasma-16D ($+10.9\%$), Levy-10D ($+24.0\%$), Levy-20D ($+9.7\%$). On Ackley-8D and Hartmann-6D, where Static KS loses, Adaptive KS is neutral (-1.4% and -0.9%, respectively). Fixed20 still loses on Hartmann (-5.0%), confirming that the decisiveness-based activation is important. Figure 2 shows effect sizes; convergence curves are given in Section E.

Robustness across initializations. Across all 55 cells (11 benchmarks $\times$ 5 initializations): 18W/2L/35N. Both losses occur on Hartmann-6D under sparse pure-matching initializations and are smaller than 2.5%. On pool initializations: 10W/0L. On the plasma problem (Shao et al., 2025): $+10.9\%$ ($p = 0.003$). The threshold τ is stable: $\tau \in \{0.25, 0.30, 0.35\}$ produces zero losses on 4 representative benchmarks (Section H).

7 Discussion and conclusion

Standard EUBO querying in PBO yields disconnected comparison graphs, causing rank deficiency in the likelihood Hessian and ill-conditioning irrespective of the inference method. KappaSharp addresses this by adding a prior-weighted diagonal correction during MAP estimation, with an adaptive variant that activates based on decisiveness. Across 11 benchmarks, including a plasma medicine task, Adaptive KappaSharp consistently improves performance, with gains up to +10.9%. Remaining limitations include an empirically chosen threshold and evaluation with oracle-simulated preferences. Future work will develop eigenvector-aligned (non-diagonal) corrections, extend analysis to larger batch sizes, and automate threshold selection.

DISCLOSURE OF LARGE LANGUAGE MODEL USE

In accordance with ICLR 2027 policy, we disclose the use of large language models (Claude, ChatGPT) during the preparation of this work.

Code implementation: LLMs assisted with implementing the experiment runner, ablation scripts, and result analysis code. All LLM-generated code was reviewed, tested, and verified for correctness by the authors.

Proof assistance: LLMs assisted in drafting and checking mathematical proofs. All proofs were independently verified by the authors and additionally reviewed by a separate LLM instance acting as an independent reviewer.

Writing and editing: LLMs were used to draft portions of the manuscript and to improve grammar, clarity, and phrasing. All LLM-generated text was reviewed and rewritten by the authors through multiple rounds of manual editing.

Figures: Figure 1 was created using ChatGPT image generation and manually refined by the authors.

Other: LLMs assisted with summarizing related work, formatting tables and references, and identifying relevant literature.

Not LLM-generated: The core research hypothesis (that standard PBO querying fragments the comparison graph and that a prior-weighted diagonal correction can address the resulting ill-conditioning), the experimental methodology, and the design of the KappaSharp correction were conceived by the authors. No synthetic datasets were generated by LLMs.

The authors have reviewed all AI-assisted work and take full responsibility for the final content of this submission, including all text, claims, and artifacts.

ETHICS STATEMENT

This work proposes a method to improve the conditioning of Gaussian process surrogates in preferential Bayesian optimization. All experiments use synthetic benchmark functions or publicly available test problems with simulated (oracle) comparisons. No human subjects were involved in the experimental evaluation. The plasma medicine benchmark uses a physics simulator, not real patients. We do not foresee negative societal impacts from this work.

REPRODUCIBILITY STATEMENT

All experiments use the open-source BoTorch library (Balandat et al., 2020) with default PairwiseGP and EUBO settings. The complete experimental protocol is described in Section 6: 11 benchmarks, 5 initialization conditions, 60 seeds per condition, paired same-node execution with common random numbers. Hyperparameters for KappaSharp ($\alpha = 0.1$, $\tau = 0.30$, $\alpha_{\text{ema}} = 0.2$, $t_{\min} = 8$) are fixed across all benchmarks and reported in Algorithm 1. The η^* calibration uses a grid search described in Section A. Benchmark descriptions and additional hyperparameters are in Section G. Code will be released upon publication.

REFERENCES

Raul Astudillo, Jerry Lin, Eytan Bakshy, and Peter I. Frazier. qEUBO: A decision-theoretic acquisition function for preferential Bayesian optimization. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2023.

Maximilian Balandat, Brian Karrer, Daniel R. Jiang, Samuel Daulton, Benjamin Letham, Andrew Gordon Wilson, and Eytan Bakshy. BoTorch: A framework for efficient Monte-Carlo Bayesian optimization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

Alessio Benavoli, Dario Azzimonti, and Dario Piga. Preferential Bayesian optimisation with skew Gaussian processes. *arXiv preprint arXiv:2008.06677*, 2020.

Eric Brochu, Vlad M. Cora, and Nando de Freitas. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.

Wei Chu and Zoubin Ghahramani. Preference learning with Gaussian processes. In *International Conference on Machine Learning (ICML)*, 2005.

Javier González, Zhenwen Dai, Andreas Damianou, and Neil D. Lawrence. Preferential Bayesian optimization. In *International Conference on Machine Learning (ICML)*, 2017.

Marcelo Hartmann and Jarno Vanhatalo. Laplace approximation and natural gradient for Gaussian process regression with heteroscedastic student-t model. *Statistics and Computing*, 29:753–773, 2019.

Alexander Immer, Maciej Korzepa, and Matthias Bauer. Improving predictions of Bayesian neural nets via local linearization. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2021.

Xiaoye Jiang, Lek-Heng Lim, Yuan Yao, and Yinyu Ye. Statistical ranking and combinatorial Hodge theory. *Mathematical Programming*, 127(1):203–244, 2011.

James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Oriol Vinyals, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114: 3521–3526, 2017.

Pascal Kündig and Fabio Sigrist. Iterative methods for Vecchia-Laplace approximations for latent Gaussian process models. *Journal of the American Statistical Association*, 2024.

Jerry Lin, Raul Astudillo, Peter I. Frazier, and Eytan Bakshy. Preference exploration for efficient Bayesian optimization with multiple outcomes. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2022.

Petrus Mikkola, Milica Todorović, Jari Järvi, Patrick Rinke, and Samuel Kaski. Projective preferential Bayesian optimization. In *International Conference on Machine Learning (ICML)*, 2020.

Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, 2006.

Nihar B. Shah, Sivaraman Balakrishnan, Joseph K. Bradley, Abhay Parekh, Kannan Ramchandran, and Martin J. Wainwright. Estimation from pairwise comparisons: Sharp minimax bounds with topology dependence. *Journal of Machine Learning Research*, 17(58):1–47, 2016.

Ketong Shao, Ankush Chakrabarty, Ali Mesbah, and Diego Romeres. Coactive preference-guided multi-objective Bayesian optimization: An application to policy learning in personalized plasma medicine. *IEEE Control Systems Letters*, 2025.

Shion Takeno et al. Towards practical preferential Bayesian optimization with skew Gaussian processes. In *International Conference on Machine Learning (ICML)*, 2023.

Maegan Tucker, Ellen Novoseller, Claudia Kann, Yanan Sui, Yisong Yue, Joel Burdick, and Aaron D. Ames. Preference-based learning for exoskeleton gait optimization. In *International Conference on Robotics and Automation (ICRA)*, 2020.

Wenjie Xu, Wenbin Wang, Yuning Jiang, Bratislav Svetozarevic, and Colin N. Jones. Principled preferential Bayesian optimization. In *International Conference on Machine Learning (ICML)*, 2024.

A PROOFS

A.1 PROOF OF PROPOSITION 1 (H_{lik} IS A WEIGHTED GRAPH LAPLACIAN)

Proof. For comparison m, the probit model (Equation (2)) gives $P(x_m^+ \succ x_m^-) = \Phi(s_m)$ with $s_m = \mathbf{C}_m^\top \mathbf{f}/(\sigma\sqrt{2})$, so the log-likelihood contribution is $\mathcal{L}_m(\mathbf{f}) = \log \Phi(s_m)$. The gradient is

$$\nabla_{\mathbf{f}} \mathcal{L}_m = \frac{1}{\sigma\sqrt{2}} \psi(s_m) \mathbf{C}_m, \quad \text{where } \psi(z) = \frac{\phi(z)}{\Phi(z)} \tag{14}$$

is the inverse Mills ratio. Differentiating again:

$$\nabla_{\mathbf{f}}^2 \mathcal{L}_m = -\frac{1}{2\sigma^2} \psi(s_m)\big(\psi(s_m) + s_m\big) \mathbf{C}_m \mathbf{C}_m^\top = -\omega_m \mathbf{C}_m \mathbf{C}_m^\top. \tag{15}$$

Positivity of ω_m. For all finite s_m, we need $\psi(s_m)(\psi(s_m) + s_m) > 0$. Since $\psi(z) = \phi(z)/\Phi(z) > 0$ for all finite z, it suffices to show $\psi(z) + z > 0$. This follows from the classical Mills ratio inequality: $\psi(z) > -z$ for all z (equivalently, $\phi(z)/\Phi(z) > -z$, which holds because $\Phi(z) > 0$ and $\phi(z) > 0$, and for $z \geq 0$ it is immediate; for $z < 0$, the bound $\Phi(z) < \phi(z)/|z|$ gives $\psi(z) > |z| = -z$).

Summing over all M comparisons yields $H_{\text{lik}} = -\nabla^2 \log p(\mathcal{D} \mid \mathbf{f}) = \sum_{m=1}^M \omega_m \mathbf{C}_m \mathbf{C}_m^\top = \mathbf{C}^\top \mathbf{W} \mathbf{C}$.

Rank. Since $\mathbf{W}$ is positive diagonal (all $\omega_m > 0$), we have $\ker(\mathbf{C}^\top \mathbf{W} \mathbf{C}) = \ker(\mathbf{W}^{1/2}\mathbf{C}) = \ker(\mathbf{C})$. The kernel of the incidence matrix $\mathbf{C}$ of a (multi)graph on N vertices with N_{cc} connected components is spanned by the N_{cc} component-indicator vectors, so $\operatorname{rank}(\mathbf{C}^\top \mathbf{W} \mathbf{C}) = N - N_{\text{cc}}$. □

A.2 PROOF OF PROPOSITION 2

Proof. Let $\mathbf{Q} = [\mathbf{Q}_0 \ \mathbf{Q}_1]$ be orthogonal with $\mathbf{Q}_0$ spanning $\ker(\mathbf{C})$ and $\mathbf{Q}_1$ spanning $\operatorname{range}(\mathbf{C}^\top)$. Write $\mathbf{f} = \mathbf{Q}_0 \xi + \mathbf{Q}_1 \zeta$.

(a) Since $\mathbf{C}\mathbf{Q}_0 = \mathbf{0}$, every comparison score $\mathbf{C}_m^\top \mathbf{f} = \mathbf{C}_m^\top \mathbf{Q}_1 \zeta$ depends only on ζ. Hence $p(\mathcal{D} \mid \xi, \zeta) = p(\mathcal{D} \mid \zeta)$.

(b) By Bayes' rule, $p(\xi, \zeta \mid \mathcal{D}) \propto p(\mathcal{D} \mid \zeta)\, p_0(\xi, \zeta) = p_0(\xi \mid \zeta)\, p(\mathcal{D} \mid \zeta)\, p_0(\zeta)$. Normalizing gives $p(\xi, \zeta \mid \mathcal{D}) = p_0(\xi \mid \zeta)\, p(\zeta \mid \mathcal{D})$, so $p(\xi \mid \zeta, \mathcal{D}) = p_0(\xi \mid \zeta)$.

Under the Laplace approximation, write the prior precision in (ξ, ζ)-coordinates as $\mathbf{Q}^\top \mathbf{K}^{-1} \mathbf{Q} = \left[\begin{smallmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^\top & \mathbf{D}_0 \end{smallmatrix}\right]$. The likelihood Hessian contributes only to the ζ-block: $\mathbf{Q}^\top \mathbf{H} \mathbf{Q} = \left[\begin{smallmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^\top & \mathbf{D}_0 + \mathbf{L} \end{smallmatrix}\right]$ with $\mathbf{L} = \mathbf{Q}_1^\top H_{\text{lik}} \mathbf{Q}_1 \succeq 0$. The MAP satisfies the first-order condition $\mathbf{A}\hat{\xi} + \mathbf{B}\hat{\zeta} = \mathbf{0}$ (since the likelihood does not depend on ξ). The Gaussian conditional on ξ given ζ therefore has mean $\hat{\xi} - \mathbf{A}^{-1}\mathbf{B}(\zeta - \hat{\zeta}) = -\mathbf{A}^{-1}\mathbf{B}\zeta$ and covariance $\mathbf{A}^{-1}$, matching the prior conditional $p_0(\xi \mid \zeta)$.

If $\mathbf{B} = \mathbf{0}$ (block-diagonal prior), then $p_0(\xi \mid \zeta) = p_0(\xi)$, so $p_L(\xi \mid \mathcal{D}) = p_0(\xi)$ and the null-space covariance is unchanged. □

A.3 RANK INVARIANCE UNDER MAP SHIFT

Proposition 10 (Rank invariance)**.** *For every finite* $\mathbf{f}$, $\ker(H_{\text{lik}}(\mathbf{f})) = \ker(\mathbf{C})$ *and* $\operatorname{rank}(H_{\text{lik}}(\mathbf{f})) = N - N_{\text{cc}}$, *since* $\mathbf{W}(\mathbf{f}) \succ 0$.

Proof. By Proposition 1, $H_{\text{lik}}(\mathbf{f}) = \mathbf{C}^\top \mathbf{W}(\mathbf{f}) \mathbf{C}$ with $\mathbf{W}(\mathbf{f}) \succ 0$. Then $v^\top H_{\text{lik}}(\mathbf{f}) v = \|\mathbf{W}(\mathbf{f})^{1/2} \mathbf{C} v\|^2$, so $v \in \ker(H_{\text{lik}}(\mathbf{f}))$ iff $\mathbf{C}v = 0$. Hence $\ker(H_{\text{lik}}(\mathbf{f})) = \ker(\mathbf{C})$ for all finite $\mathbf{f}$, and $\operatorname{rank}(H_{\text{lik}}(\mathbf{f})) = N - N_{\text{cc}}$ by rank-nullity. □

A.4 PROOF OF PROPOSITION 4 (CORRECTION PROPERTIES)

Proof. Let $a_i = (\mathbf{K}^{-1})_{ii} > 0$ and $r_i(\eta) = \eta^2/(\eta + a_i)$ for $\eta > 0$.

(i) Positivity. $\eta > 0$ and $\eta + a_i > 0$ immediately give $r_i > 0$.

(ii) Monotonicity in η. $\frac{\partial r_i}{\partial \eta} = \frac{\eta(\eta+2a_i)}{(\eta+a_i)^2}$. The numerator is positive (both factors positive for $\eta > 0$, $a_i > 0$), and the denominator is positive.

(iii) Decreasing in a_i. $\frac{\partial r_i}{\partial a_i} = -\frac{\eta^2}{(\eta+a_i)^2} < 0$.

(iv) Bounded. $r_i = \eta - \frac{\eta a_i}{\eta+a_i} < \eta$ since $\frac{\eta a_i}{\eta+a_i} > 0$.

(v)–(vi) Limits. As $a_i \to 0$: $r_i/\eta = \eta/(\eta + a_i) \to 1$. As $a_i \to \infty$: $r_i/\eta = \eta/(\eta + a_i) \to 0$.

Variational form. The objective $q(r) = \frac{a_i}{2}r^2 + \frac{\eta}{2}(r-\eta)^2$ is strictly convex with $q''(r) = a_i + \eta > 0$. Setting $q'(r) = (a_i + \eta)r - \eta^2 = 0$ gives $r^* = \eta^2/(\eta + a_i) = r_i$. □

A.5 Proof of Proposition 5 (κ-Calibration is Well-Posed)

Proof. Continuity and PSD. Each $r_i(\eta)$ is a continuous, non-negative function of $\eta \geq 0$ (with $r_i(0) = 0$), so $\mathbf{R}(\eta) = \mathrm{diag}(r_1, \ldots, r_N)$ is continuous and PSD.

Loewner monotonicity. For $\eta' > \eta \geq 0$: each $r_i(\eta') > r_i(\eta)$ by property (ii), so $\mathbf{R}(\eta') - \mathbf{R}(\eta)$ is a positive diagonal matrix, hence $\mathbf{R}(\eta') \succ \mathbf{R}(\eta)$ in the Loewner order.

Limiting behavior. As $\eta \to \infty$, $r_i(\eta)/\eta \to 1$ for all i, so $\mathbf{R}(\eta) \approx \eta\mathbf{I}$ for large η. Since $\mathbf{H}$ and $\mathbf{R}(\eta)$ are both symmetric but do not commute in general, we bound $\kappa(\mathbf{H} + \mathbf{R}(\eta))$ via Weyl's inequalities. By Weyl's inequality for eigenvalue sums of Hermitian matrices:

$$\lambda_{\max}(\mathbf{H} + \mathbf{R}) \leq \lambda_{\max}(\mathbf{H}) + \lambda_{\max}(\mathbf{R}) \leq \lambda_{\max}(\mathbf{H}) + \eta, \tag{16}$$

$$\lambda_{\min}(\mathbf{H} + \mathbf{R}) \geq \lambda_{\min}(\mathbf{H}) + \lambda_{\min}(\mathbf{R}) = \lambda_{\min}(\mathbf{H}) + \min_i r_i(\eta), \tag{17}$$

where the second line uses that $\mathbf{R}$ is diagonal with eigenvalues $r_i(\eta)$. Since $r_i(\eta) = \eta^2/(\eta + a_i) = \eta(1 - a_i/(\eta + a_i))$, we have $\min_i r_i(\eta) = \eta(1 - a_{\max}/(\eta + a_{\max}))$ where $a_{\max} = \max_i a_i$. As $\eta \to \infty$, $\min_i r_i(\eta)/\eta \to 1$. Therefore:

$$\kappa(\mathbf{H} + \mathbf{R}(\eta)) \leq \frac{\lambda_{\max}(\mathbf{H}) + \eta}{\lambda_{\min}(\mathbf{H}) + \eta(1 - a_{\max}/(\eta + a_{\max}))} \to 1 \quad \text{as } \eta \to \infty. \tag{18}$$

Existence of η^*. If $\kappa(\mathbf{H}) = 1$, the target $\kappa(\mathbf{H})^{1-\alpha} = 1$ is already met at $\eta^* = 0$. Otherwise, the function $g(\eta) = \kappa(\mathbf{H} + \mathbf{R}(\eta))$ is continuous, $g(0) = \kappa(\mathbf{H}) > 1$, and $g(\eta) \to 1 < \kappa(\mathbf{H})^{1-\alpha}$ for $\alpha \in (0, 1)$. By the intermediate value theorem, $\{\eta \geq 0 : g(\eta) \leq \kappa(\mathbf{H})^{1-\alpha}\}$ is nonempty. As a preimage of a closed set under a continuous function, it is closed; its intersection with $[0, M]$ for sufficiently large M is compact, yielding the minimum η^*.

Remark 11 (Implementation). We do not claim monotonicity of $g(\eta)$ in η; non-commutativity of $\mathbf{H}$ and $\mathbf{R}$ prevents this in general. The implementation searches a log-spaced grid of 60 points from 10^{-6} to 10^8, identifies the first grid interval where the threshold is crossed, and refines with 30 bisection steps within that interval. The existence result guarantees a valid η^* exists on $[0, \infty)$; in all experiments, a feasible point was found.

□

A.6 Proof of Proposition 6 (Controlled MAP Shift)

Proof. Let $\hat{\mathbf{f}}_0$ satisfy $\nabla J_0(\hat{\mathbf{f}}_0) = 0$ (baseline MAP) and $\hat{\mathbf{f}}_\eta$ satisfy $\nabla J_0(\hat{\mathbf{f}}_\eta) + \mathbf{R}\hat{\mathbf{f}}_\eta = 0$ (augmented MAP). Define $\Delta = \hat{\mathbf{f}}_\eta - \hat{\mathbf{f}}_0$.

By the fundamental theorem of calculus:

$$\nabla J_0(\hat{\mathbf{f}}_\eta) - \nabla J_0(\hat{\mathbf{f}}_0) = \left(\int_0^1 \nabla^2 J_0(\hat{\mathbf{f}}_0 + t\Delta)\, dt \right) \Delta = \bar{\mathbf{H}}\Delta. \tag{19}$$

From the augmented optimality: $\nabla J_0(\hat{\mathbf{f}}_\eta) = -\mathbf{R}\hat{\mathbf{f}}_\eta = -\mathbf{R}\hat{\mathbf{f}}_0 - \mathbf{R}\Delta$. Substituting:

$$\bar{\mathbf{H}}\Delta = -\mathbf{R}\hat{\mathbf{f}}_0 - \mathbf{R}\Delta \implies (\bar{\mathbf{H}} + \mathbf{R})\Delta = -\mathbf{R}\hat{\mathbf{f}}_0. \tag{20}$$

Since $\bar{\mathbf{H}} \succeq \mathbf{K}^{-1} \succ 0$ and $\mathbf{R} \succeq 0$, the matrix $\bar{\mathbf{H}} + \mathbf{R}$ is invertible. Setting $\mathbf{M} = \bar{\mathbf{H}}^{-1/2}\mathbf{R}\bar{\mathbf{H}}^{-1/2}$:

$$\bar{\mathbf{H}}^{1/2}\Delta = -(\mathbf{I} + \mathbf{M})^{-1}\mathbf{M}\bar{\mathbf{H}}^{1/2}(\hat{\mathbf{f}}_0). \tag{21}$$

Taking norms: $\|\Delta\|_{\bar{\mathbf{H}}} = \|\bar{\mathbf{H}}^{1/2}\Delta\| \leq \|(\mathbf{I}+\mathbf{M})^{-1}\mathbf{M}\| \cdot \left\|\hat{\mathbf{f}}_0\right\|_{\bar{\mathbf{H}}}$. Since $\mathbf{M} \succeq 0$ with largest eigenvalue $\mu_{\max}$, the operator norm $\|(\mathbf{I} + \mathbf{M})^{-1}\mathbf{M}\| = \mu_{\max}/(1 + \mu_{\max})$. □

A.7 Proof of Proposition 8 (Full Augmentation Contracts EUBO)

Proof. Both posteriors share the MAP $\hat{\mathbf{f}}_\eta$, so predictive means coincide: $\delta_{ab} = c_{ab}^\top \hat{\mathbf{f}}_\eta$ for contrast $c_{ab} = e_a - e_b$.

Variance ordering. Since $\mathbf{R} \succeq 0, \mathbf{H}_\eta + \mathbf{R} \succeq \mathbf{H}_\eta$, so $(\mathbf{H}_\eta + \mathbf{R})^{-1} \preceq \mathbf{H}_\eta^{-1}$ (inversion reverses Loewner order on the PD cone). Therefore $s^2_{\mathrm{aug}}(ab) = c_{ab}^\top(\mathbf{H}_\eta + \mathbf{R})^{-1}c_{ab} \leq c_{ab}^\top \mathbf{H}_\eta^{-1} c_{ab} = s^2_{\mathrm{MO}}(ab)$.

EUBO monotonicity. For q=2 EUBO with $\mathrm{EUBO}(\delta, s) = \delta\Phi(\delta/s) + s\phi(\delta/s)$ (Equation (6)), differentiate at fixed δ: $\partial\,\mathrm{EUBO}/\partial s = \phi(\delta/s) > 0$ for $s > 0$. (The $z^2\phi(z)$ terms from the chain rule cancel exactly.) Hence $s_{\mathrm{MO}} \geq s_{\mathrm{aug}}$ implies $\mathrm{EUBO}_{\mathrm{MO}} \geq \mathrm{EUBO}_{\mathrm{aug}}$.

Distortion bound. By the mean value theorem, $\mathrm{EUBO}_{\mathrm{MO}} - \mathrm{EUBO}_{\mathrm{aug}} = \phi(\delta/\tilde{s})(s_{\mathrm{MO}} - s_{\mathrm{aug}})$ for some $\tilde{s} \in [s_{\mathrm{aug}}, s_{\mathrm{MO}}]$. Since $\phi \leq 1/\sqrt{2\pi}$, we get $\mathrm{EUBO}_{\mathrm{MO}} - \mathrm{EUBO}_{\mathrm{aug}} \leq (s_{\mathrm{MO}} - s_{\mathrm{aug}})/\sqrt{2\pi}$.

Spectral control. Let $\mathbf{M} = \mathbf{H}_\eta^{-1/2}\mathbf{R}\mathbf{H}_\eta^{-1/2}$ with eigenvalues λ_i. Then $s^2_{\mathrm{MO}}/s^2_{\mathrm{aug}} \leq 1 + \lambda_{\max}(\mathbf{M})$, so the distortion is controlled by the spectrum of $\mathbf{H}_\eta^{-1/2}\mathbf{R}\mathbf{H}_\eta^{-1/2}$.

Extension to fresh query points. Let $x_a, x_b \in \mathcal{X}$ be arbitrary query points (possibly $x_a, x_b \notin X$), and define $r_{ab} := k(X, x_a) - k(X, x_b) \in \mathbb{R}^N$ and $\nu_{ab} := k(x_a, x_a) + k(x_b, x_b) - 2k(x_a, x_b) - r_{ab}^\top \mathbf{K}^{-1} r_{ab} \geq 0$ (nonnegative by the Schur complement of the joint GP prior). Under a Laplace posterior $\mathcal{N}(\hat{\mathbf{f}}_\eta, \Sigma)$ on the training latents, the GP conditional gives $f(x_a) - f(x_b) \mid \mathbf{f} \sim \mathcal{N}(r_{ab}^\top \mathbf{K}^{-1}\mathbf{f}, \nu_{ab})$, so by total expectation and variance:

$$\delta_{ab} = r_{ab}^\top \mathbf{K}^{-1}\hat{\mathbf{f}}_\eta, \quad s^2_\Sigma(ab) = \nu_{ab} + r_{ab}^\top \mathbf{K}^{-1}\Sigma\,\mathbf{K}^{-1} r_{ab}. \tag{22}$$

For MAP-only, $\Sigma_{\mathrm{MO}} = \mathbf{H}_\eta^{-1}$; for full augmentation, $\Sigma_{\mathrm{aug}} = (\mathbf{H}_\eta + \mathbf{R})^{-1}$. Since $(\mathbf{H}_\eta + \mathbf{R})^{-1} \preceq \mathbf{H}_\eta^{-1}$ and congruence preserves Loewner order, $r_{ab}^\top \mathbf{K}^{-1}(\mathbf{H}_\eta + \mathbf{R})^{-1}\mathbf{K}^{-1} r_{ab} \leq r_{ab}^\top \mathbf{K}^{-1}\mathbf{H}_\eta^{-1}\mathbf{K}^{-1} r_{ab}$, giving $s^2_{\mathrm{aug}}(ab) \leq s^2_{\mathrm{MO}}(ab)$. The mean δ_{ab} is identical (shared MAP). EUBO monotonicity then gives $\mathrm{EUBO}_{\mathrm{aug}}(a, b) \leq \mathrm{EUBO}_{\mathrm{MO}}(a, b)$ for all pairs, including fresh query points. When $x_a, x_b \in X$, $r_{ab} = \mathbf{K}(e_a - e_b)$ and $\nu_{ab} = 0$, recovering the in-sample case. □

A.8 Proof of Theorem 7 (Pseudo-Confidence in Null Directions)

Proof. In (ξ, ζ)-coordinates, the MAP-only precision is $\mathbf{Q}^\top \mathbf{H}_\eta \mathbf{Q} = \left[\begin{smallmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^\top & \mathbf{D}_\eta \end{smallmatrix}\right]$ (with $\mathbf{A} = \mathbf{Q}_0^\top \mathbf{K}^{-1}\mathbf{Q}_0$, since $\mathbf{Q}_0^\top H_{\mathrm{lik}}\mathbf{Q}_0 = \mathbf{0}$), and the augmented precision adds $\mathbf{Q}^\top \mathbf{R}\mathbf{Q} = \left[\begin{smallmatrix} \mathbf{R}_{00} & \mathbf{R}_{01} \\ \mathbf{R}_{01}^\top & \mathbf{R}_{11} \end{smallmatrix}\right]$. The Gaussian conditional covariance of ξ given ζ equals the inverse of the (ξ, ξ)-block of the precision. For MAP-only this is $\mathbf{A}^{-1}$; for full augmentation it is $(\mathbf{A} + \mathbf{R}_{00})^{-1}$. Since $\mathbf{R} \succ 0$ and $\mathbf{Q}_0$ has full column rank, $\mathbf{R}_{00} = \mathbf{Q}_0^\top \mathbf{R}\mathbf{Q}_0 \succ 0$, so $(\mathbf{A} + \mathbf{R}_{00})^{-1} \prec \mathbf{A}^{-1}$. □

A.9 Proof of Theorem 9 (Argmax Stability)

Proof. Let $(a^*_{\mathrm{aug}}, b^*_{\mathrm{aug}})$ maximize $\mathrm{EUBO}_{\mathrm{aug}}$. Telescoping gives $\mathrm{EUBO}_{\mathrm{MO}}(a^*_{\mathrm{MO}}, b^*_{\mathrm{MO}}) - \mathrm{EUBO}_{\mathrm{MO}}(a^*_{\mathrm{aug}}, b^*_{\mathrm{aug}}) \leq 2\sup_{a,b} |\mathrm{EUBO}_{\mathrm{MO}}(a, b) - \mathrm{EUBO}_{\mathrm{aug}}(a, b)|$. By the mean value theorem and $\partial\,\mathrm{EUBO}/\partial s = \phi(\delta/s) \leq 1/\sqrt{2\pi}$, each pointwise gap is at most $(s_{\mathrm{MO}} - s_{\mathrm{aug}})/\sqrt{2\pi}$. For each pair (a, b), let b_{ab} denote the contrast vector such that $s^2_{\mathrm{MO}}(a, b) = \nu_{ab} + b_{ab}^\top \mathbf{H}_\eta^{-1} b_{ab}$ (with $\nu_{ab} \geq 0$ the Schur-complement term from Proposition 8). Then $s^2_{\mathrm{aug}} \geq s^2_{\mathrm{MO}}/(1 + \lambda_1)$, so $s_{\mathrm{MO}} - s_{\mathrm{aug}} \leq s_{\mathrm{MO}}\lambda_1/2$. Combining: the gap is at most $\sup_{a,b} s_{\mathrm{MO}}(a, b)\lambda_1/\sqrt{2\pi}$. The supremum is finite when the kernel k is continuous and $\mathcal{X}$ is compact: by the Cauchy–Schwarz inequality for positive-definite kernels, $|k(x, y)| \leq \sqrt{k(x, x)k(y, y)}$, so $s^2_{\mathrm{MO}}(a, b) \leq 4\sup_x k(x, x)$ for all (a, b). Since $\lambda_1 \leq \|\mathbf{R}\|/\lambda_{\min}(\mathbf{H}_\eta) < \eta/\lambda_{\min}(\mathbf{H}_\eta)$, the gap is $O(\eta/\lambda_{\min}(\mathbf{H}_\eta))$. □

Table 2: Mechanism ablation: gain (%) over baseline B (pool_n6_k15, 60 seeds). Static KS: full prior-gated diagonal correction. Jitter: uniform $R{=}cI$ (no prior gating). Diagonal: uniform diagonal regularization ($\hat{R}{=}\eta I$). MetricOnly / CovOnly: partial mechanisms. *: $p{<}0.05$, **: $p{<}0.01$ (paired t-test vs. B).

Benchmark	d	Static KS	Jitter	Diagonal	MetricOnly	CovOnly
DTLZ2-16D	16	**+24.5****	-22.6**	**+30.7****	+4.3	+7.0
DTLZ2-8D	8	**+23.2***	-47.8**	+1.0	+13.3	-14.9
DTLZ2-20D	20	**+17.8****	-21.0**	**+18.9****	+4.7	+2.4
Plasma-16D	16	**+9.5****	+0.4	**+8.9***	+2.3	-2.9
Levy-10D	10	**+18.7***	-10.3	+10.5	**+18.5***	**+20.9***
Ackley-8D	8	-15.6*	-45.0**	-22.8**	+2.2	-15.4
Hartmann-6D	6	-4.8**	-7.0**	-6.7**	+0.2	-2.1
Vehicle-5D	5	-0.0	-0.0	-0.0	+0.0	-0.0
W / L / N		5/2/1	0/5/3	3/2/3	1/0/7	1/0/7

A.10 Frozen-Utility Monotonicity

Proposition 12 (Frozen-utility monotonicity)**.** *Fix latent utilities* $\mathbf{f} \in \mathbb{R}^N$. *Let* $H_{\text{lik}}(\mathbf{f}) = \sum_{m=1}^{M} \omega_m(\mathbf{f})\, \mathbf{C}_m \mathbf{C}_m^\top$ *and let* $H'_{\text{lik}}(\mathbf{f})$ *denote the Hessian after removing comparison* m'*, with all weights evaluated at the same fixed* $\mathbf{f}$*. Then* $H'_{\text{lik}}(\mathbf{f}) = H_{\text{lik}}(\mathbf{f}) - \omega_{m'}(\mathbf{f})\, \mathbf{C}_{m'} \mathbf{C}_{m'}^\top \preceq H_{\text{lik}}(\mathbf{f})$. *This is a frozen-curvature statement: it does not compare Hessians obtained after refitting the MAP on the reduced dataset, since refitting changes the remaining weights.*

Proof. $\omega_{m'}(\mathbf{f})\, \mathbf{C}_{m'} \mathbf{C}_{m'}^\top \succeq 0$ since $\omega_{m'}(\mathbf{f}) > 0$, so $H_{\text{lik}}(\mathbf{f}) - H'_{\text{lik}}(\mathbf{f}) \succeq 0$. □

B Mechanism Ablation

Uniform jitter adds a fixed constant c to every diagonal entry of the posterior precision ($R{=}cI$, where c is chosen to match the trace of the KS correction); it scores 0W/5L. Diagonal regularization ($\hat{R}{=}\eta I$) uses the κ-calibrated strength η^* but applies it uniformly without the $(\mathbf{K}^{-1})_{ii}$ weighting; it scores 3W/2L, with $+30.7\%$ on DTLZ2-16D but the same losses as Static KS.

MetricOnly adds $\mathbf{R}$ to the Hessian used in the Newton step for computing the MAP, but evaluates J_0 (the original objective without $\mathbf{R}$) at each Newton iterate. Since the objective being minimized is unchanged, the MAP estimate is the same as the baseline; only the Newton search direction is affected. CovOnly keeps the standard MAP estimate $\hat{\mathbf{f}}_0$ but uses $(\mathbf{H} + \mathbf{R})^{-1}$ instead of $\mathbf{H}^{-1}$ as the Laplace covariance. Both score 1W/0L/7N, indicating that the improvement comes from shifting the MAP estimate (as in the full KappaSharp correction), not from changing the Newton search direction or tightening the covariance alone.

C Initialization Robustness

C.1 Proof of Proposition 3 (Effective Information Dimension)

Proof. Let $\tilde{\mathbf{M}} = \mathbf{K}^{1/2} H_{\text{lik}} \mathbf{K}^{1/2} \succeq 0$ with eigenvalues $\mu_1 \geq \cdots \geq \mu_N \geq 0$. Since H_{lik} has rank $N{-}N_{\text{cc}}$ (Proposition 1), $\tilde{\mathbf{M}}$ has exactly N_{cc} zero eigenvalues.

(a) $d_{\text{eff}} = \sum_{i=1}^{N} \mu_i/(1+\mu_i)$. Each term satisfies $0 \leq \mu_i/(1+\mu_i) < 1$, with equality to 0 iff $\mu_i = 0$. Since exactly N_{cc} eigenvalues are zero, at most $N{-}N_{\text{cc}}$ terms are positive, giving $d_{\text{eff}} \leq N{-}N_{\text{cc}}$. Equality requires all nonzero $\mu_i \to \infty$.

(b) Under the fresh-pair assumption with $N_{\text{cc},0}$ initial connected components: each of the T queries adds two unseen points and one new component, so $N_{\text{cc}} = N_{\text{cc},0} + T$ and $N = 2T + N_0$. Therefore $\text{rank}(H_{\text{lik}}) = N - N_{\text{cc}} = N_0 + T - N_{\text{cc},0}$ and $d_{\text{eff}} \leq N_0 + T - N_{\text{cc},0}$. Hence $d_{\text{eff}}/N \leq (N_0 + T - N_{\text{cc},0})/(2T + N_0) \to 1/2$ as $T \to \infty$ (with $N_0, N_{\text{cc},0}$ fixed).

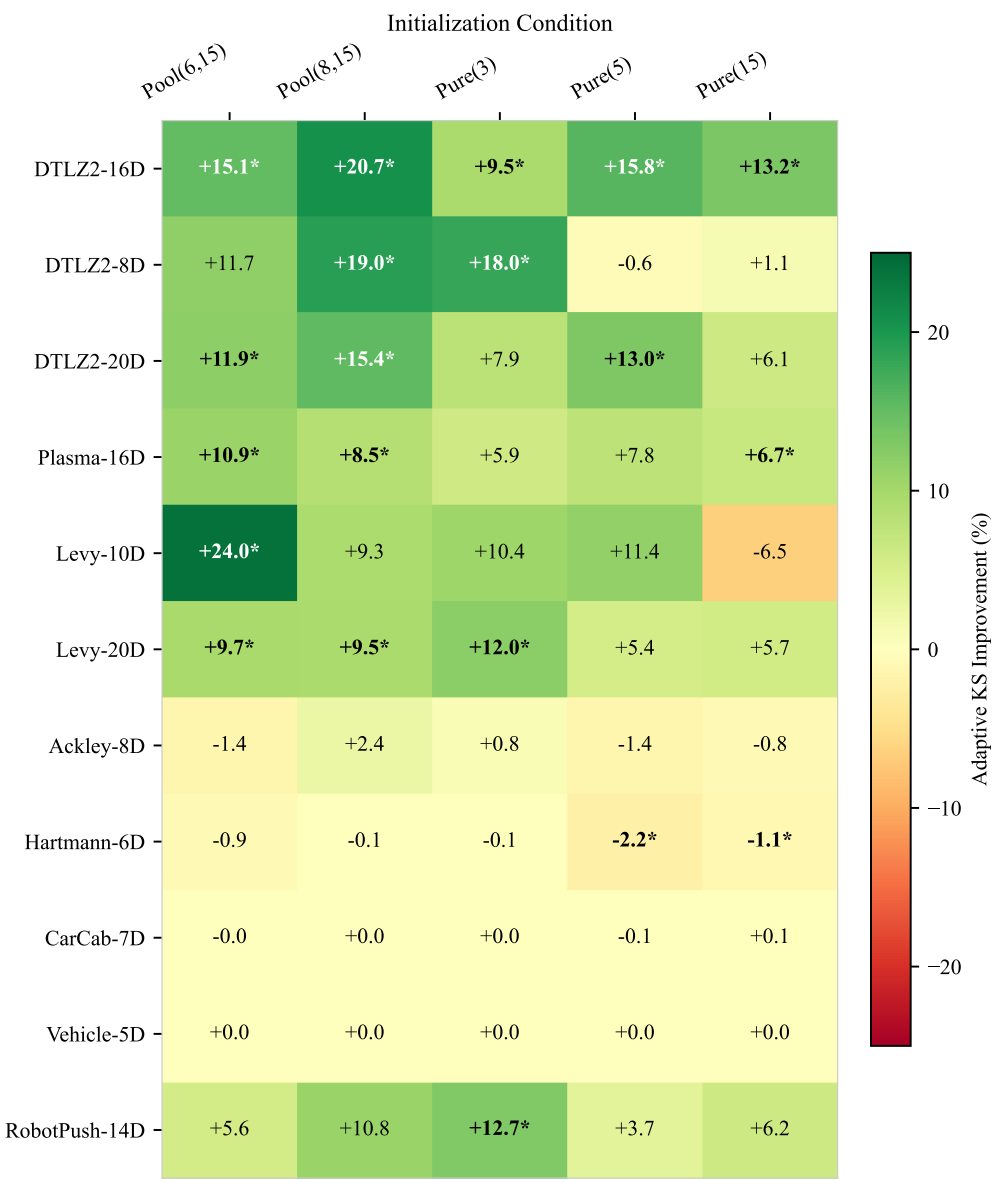

Figure 3: Adaptive KS improvement (%) over baseline for each of the 55 benchmark–initialization cells. Each cell shows the mean improvement over 60 seeds. Bold with asterisk (*): $p < 0.05$ (paired t-test). Green = improvement, red = regression. The two losses (Hartmann-6D under pure_k5 and pure_k15) are both below 2.5%.

(c) Connected (N_{cc}=1): $d_{\text{eff}} \leq N-1$. The rank ceiling is lifted to $N-1$, though the realized value depends on the magnitude of comparison weights (a connected graph with vanishing edge weights can still have $d_{\text{eff}} \ll N-1$).

(d) Let $\mathbf{B} = \mathbf{K}^{1/2}\mathbf{R}\mathbf{K}^{1/2} \succeq \beta\mathbf{I}$ with $\beta > 0$. Define $g(\mathbf{X}) = \operatorname{tr}\big(\mathbf{X}(\mathbf{I}+\mathbf{X})^{-1}\big) = N - \operatorname{tr}\big((\mathbf{I}+\mathbf{X})^{-1}\big)$. Since g is monotone increasing under the Loewner order (adding PSD mass decreases $(\mathbf{I}+\mathbf{X})^{-1}$), and $\tilde{\mathbf{M}} + \mathbf{B} \succeq \tilde{\mathbf{M}} + \beta\mathbf{I}$, we have $d_{\text{eff}}^{\text{aug}} = g(\tilde{\mathbf{M}} + \mathbf{B}) \geq g(\tilde{\mathbf{M}} + \beta\mathbf{I})$. Since $\beta\mathbf{I}$ commutes with $\tilde{\mathbf{M}}$:

$$g(\tilde{\mathbf{M}} + \beta\mathbf{I}) = \sum_{i=1}^{N} \frac{\mu_i + \beta}{1 + \mu_i + \beta} = d_{\text{eff}} + \sum_{i=1}^{N} \frac{\beta}{(1+\mu_i)(1+\mu_i+\beta)}. \tag{23}$$

Each term in the sum is nonnegative, and the N_{cc} terms with $\mu_i = 0$ each contribute $\beta/(1+\beta)$, giving the bound $d_{\text{eff}}^{\text{aug}} \geq d_{\text{eff}} + N_{\text{cc}}\beta/(1+\beta)$. □

D THRESHOLD INTERPRETATION

Proposition 13 (Per-step decisiveness threshold)**.** *Since $p_t = \Phi(|\cdot|) \in [1/2, 1)$ by construction, the mapping $p \mapsto s_{\text{exp}} = 1 - \mathcal{H}(p)/\log 2$ is strictly increasing on $[1/2, 1)$. Hence $s_{\text{exp},t} \geq \tau$ implies $p_t \geq p(\tau)$. At $\tau = 0.30$: $p_t \geq 0.81$.*

Proof. $\mathcal{H}(p)/\log 2$ is maximized at $p = 0.5$ and strictly decreasing on $(0.5, 1)$. Solving $1 - \mathcal{H}(p)/\log 2 = 0.30$ yields $p \approx 0.8107$. □

On multi-objective benchmarks where the disconnected-graph problem is severe, the model quickly becomes confident about queried comparisons and $\bar{s}_t$ rises above τ, activating the correction. On smooth single-objective basins (e.g., Ackley), the model remains uncertain during exploration and $\bar{s}_t$ stays low, keeping the correction off.

E CONVERGENCE CURVES

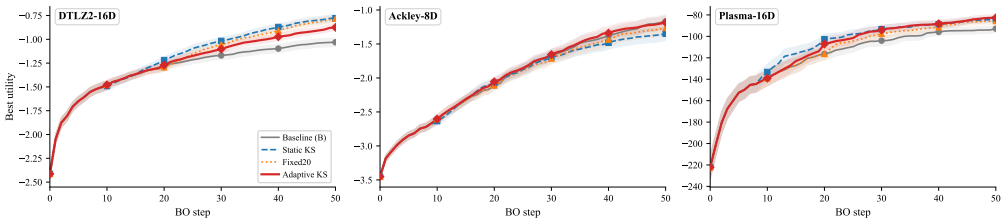

Figure 4: Convergence (mean ± s.e., 60 seeds). Left: DTLZ2-16D. Center: Ackley-8D. Right: Plasma-16D (+10.9%).

F ACTIVATION BEHAVIOR

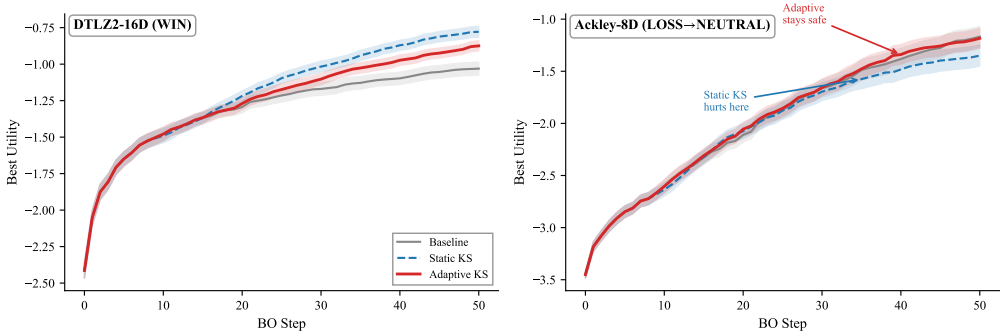

Figure 5: Activation behavior on WIN (DTLZ2-16D) vs. LOSS→NEUTRAL (Ackley-8D) benchmarks. On DTLZ2-16D, the activation rule enables the correction frequently (median ≈ 40% of steps); on Ackley-8D, low model decisiveness keeps the correction largely off (median ≈ 12%), preventing the MAP shift that causes Static KS's loss.

G Additional Experimental Details

G.1 Benchmark Descriptions

DTLZ2 (8D, 16D, 20D). Standard multi-objective test function (González et al., 2017) with d-dimensional input mapped to a 4-dimensional outcome space, then to scalar utility via a piecewise-linear aggregation. The Pareto front is a unit hypersphere in outcome space.

Plasma-16D. Real industrial application: calibrating a 16-parameter model predictive controller for atmospheric plasma surface treatment (Shao et al., 2025). Each function evaluation solves an MPC optimization with physics simulation. The utility combines temperature uniformity and treatment speed, assessed by human inspection of temperature distributions.

Ackley-8D. Standard global optimization benchmark on $[-1, 1]^8$ with many shallow local optima surrounding a single global minimum. The smooth basin structure provides strong local curvature, making the Laplace approximation already well-conditioned.

Hartmann-6D. Standard 6-dimensional test function with 6 local minima. Like Ackley, the smooth landscape provides sufficient local curvature for the baseline.

Levy-10D, Levy-20D. Multi-modal test functions with increasingly complex landscape at higher dimensions. The 20D variant tests scalability.

CarCab-7D, Vehicle-5D, RobotPush-14D. CarCab: 7-dimensional vehicle cabin design. Vehicle: 5-dimensional crash safety optimization. RobotPush: 14-dimensional robotic pushing task simulated with Box2D physics, featuring contact discontinuities.

G.2 Hyperparameters

All experiments use BoTorch 0.6.0, GPyTorch 1.6.0, PyTorch 1.11.0. PairwiseGP uses default RBF kernel with automatic relevance determination (ARD). EUBO is optimized with 4 restarts and 32 raw samples. The precision correction activates at step 8 with $\alpha = 0.1$. The η^* solve uses a grid search on 60 log-spaced points from 10^{-6} to 10^8.

Init tables are generated with hash-based independent RNG per (benchmark, seed, stream) using `gen_init_tables_v3.py`. Pool initializations use nested point sets ($n_6 \subset n_8$) with independently shuffled comparison orderings.

Activation step. The correction is eligible from step 8 by default. In a preliminary study using Static KS on DTLZ2-16D (30 seeds, separate from the main 60-seed runs), we tested eligibility at steps 5, 8, 10, and 12; final improvement ranged from $+19\%$ to $+24\%$, indicating low sensitivity to the exact activation step. Earlier eligibility (step 5) occasionally produced unstable η^* estimates; later eligibility (step 12) delayed the benefit without improving it. All main-table results use step 8 with 60 seeds.

Compute. Experiments were run on a university HPC cluster (AMD EPYC 7543 and Intel Xeon Platinum CPUs, 64-core nodes). Total compute: approximately 5,000 CPU-hours across $16{,}500$ BO runs in the main grid plus ablations and sensitivity studies. All 60-seed conditions use paired same-node execution with common random numbers (CRN).

H Threshold Sensitivity

The threshold $\tau = 0.30$ has a per-step interpretation: when $s_{\exp,t} \geq 0.30$, the model's confidence in its own prediction satisfies $p_t \geq 0.81$ (Proposition 13). The implementation triggers on the EMA $\bar{s}_t$ rather than the raw score; since $s_{\exp}$ is nonlinear in p, $\bar{s}_t \geq \tau$ does not imply an EMA-weighted average $p \geq 0.81$, but it does indicate that recent steps have exhibited high decisiveness on average. Lower τ (more aggressive) increases activation rate and gains on WIN benchmarks. All three thresholds produce zero statistically significant losses.

Table 3: Threshold sensitivity: Adaptive KS gain (%) over baseline for $\tau \in \{0.25, 0.30, 0.35\}$ on 4 representative benchmarks (pool_n6_k15, 60 seeds). All three thresholds produce zero losses. *: $p<0.05$, **: $p<0.01$ (paired t-test).

Benchmark	$\tau = 0.25$	$\tau = 0.30$	$\tau = 0.35$
DTLZ2-16D	+20.2**	+15.1**	+20.1**
Levy-10D	+25.0**	+24.0**	+22.4**
Ackley-8D	-2.3	-1.4	+4.3
Plasma-16D	+9.9**	+10.9**	+7.0*
Losses	0	0	0

I Calibration Parameter α Sensitivity

Table 4: Static KS gain (%) over baseline for $\alpha \in \{0.05, 0.1, 0.2\}$ on 4 benchmarks (pool_n6_k15). Larger α targets a lower condition number, producing stronger augmentation. $\alpha = 0.1$ is a reasonable operating point; the tradeoff is not monotone (e.g., Ackley at α=0.05 is worse than at α=0.1). *: $p<0.05$, **: $p<0.01$.

Benchmark	$\alpha = 0.05$	$\alpha = 0.1$	$\alpha = 0.2$
DTLZ2-16D	+18.8**	+24.5**	+31.8**
Plasma-16D	+6.7	+9.5**	+17.4**
Ackley-8D	-31.4^{*}	-15.6^{*}	-51.0^{**}
Hartmann-6D	-1.0	-4.8^{**}	-11.8^{**}

J Connected Querying Full Results

Table 5: Connected querying (Xu et al., 2024) detailed results (pool_n6_k15, 60 seeds). One candidate per query is fixed to the previous query point ($x'_t = x_{t-1}$); EUBO is retained as the acquisition function for fair comparison. *: $p < 0.05$, **: $p < 0.01$.

Benchmark	d	Conn. Query Gain (%)	Verdict
DTLZ2-16D	16	$+1.5$	NEUTRAL
DTLZ2-8D	8	-93.9^{**}	LOSS
DTLZ2-20D	20	$+6.2$	NEUTRAL
Plasma-16D	16	-11.5^{*}	LOSS
Levy-10D	10	-21.7^{*}	LOSS
Levy-20D	20	-8.7	NEUTRAL
Ackley-8D	8	-61.9^{**}	LOSS
Hartmann-6D	6	-11.8^{**}	LOSS
CarCab-7D	7	-0.2	NEUTRAL
Vehicle-5D	5	-0.0^{*}	LOSS
RobotPush-14D	14	-24.8^{**}	LOSS
W / L / N		0 / 7 / 4	

Connected querying enforces a connected comparison graph by fixing one candidate per query to the previous query point ($x'_t = x_{t-1}$). Note that the original algorithm of Xu et al. (2024) uses a different acquisition function based on optimistic likelihood-ratio confidence sets; here we adopt only their querying strategy and retain EUBO as the acquisition function so that the comparison isolates the effect of graph connectivity. Fixing one candidate to the previous point limits exploration, resulting in losses on 7 of 11 benchmarks, including -93.9% on DTLZ2-8D and -61.9% on Ackley-8D.

K Non-Laplace Inference Baselines

EP-EI and Gibbs-EI results on 4 representative benchmarks (pool_n6_k15, 60 seeds):

Benchmark	EP-EI Gain	Verdict	Gibbs-EI Gain	Verdict
DTLZ2-16D	$-75.2\%^{**}$	LOSS	$+9.6\%$	NEUTRAL
Ackley-8D	$-130.2\%^{**}$	LOSS	$-52.5\%^{**}$	LOSS
Plasma-16D	$-73.1\%^{**}$	LOSS	$-40.3\%^{**}$	LOSS
Levy-10D	$-82.9\%^{**}$	LOSS	$-22.0\%^{*}$	LOSS

Table 6: Non-Laplace inference baselines (diagnostic). Both use EI acquisition, not EUBO, which contributes to the poor performance. *: $p < 0.05$, **: $p < 0.01$.

L Wall-Clock Overhead

Table 7: Wall-clock time (seconds) per 50-step BO run. KappaSharp adds negligible overhead and is often *faster* due to improved Hessian conditioning.

Benchmark	Baseline (s)	KappaSharp (s)	Overhead
DTLZ2-16D	120±42	148±43	+23%
DTLZ2-8D	126±46	129±43	+2%
Plasma-16D	272±81	271±83	-0%
Levy-10D	156±47	146±38	-7%
Ackley-8D	126±44	128±43	+2%
Hartmann-6D	118±50	102±36	-13%
Vehicle-5D	48±10	54±9	+11%
RobotPush-14D	831±84	817±80	-2%

The correction adds per-step overhead when active. On 4 of 8 benchmarks, the method is actually *faster* than baseline (-2% to -13% wall-clock) because the better-conditioned Hessian reduces Newton iterations; on others, the extra MAP solve adds up to $+23\%$ wall-clock time.

M Full Per-Initialization Numeric Results

The heatmap in Figure 3 summarizes Adaptive KS's improvement across all 55 benchmark–initialization conditions. Numeric values for each cell are available in the figure. Key observations:

- **Pool initializations** (n6_k15, n8_k15): 10W/0L/12N across all benchmarks. The nested pool structure ($n_6 \subset n_8$) provides a dense initial comparison graph that combines well with the activation rule.
- **Pure matching initializations** (k3, k5, k15): 8W/2L/23N. Both losses occur on Hartmann-6D (pure_k5: -2.2%, $p < 0.05$; pure_k15: -1.1%, $p < 0.05$). These are the smallest significant effects in the entire study, and both occur on a benchmark where the baseline is already well-conditioned.
- **High-dimensional benchmarks** (DTLZ2-16D, DTLZ2-20D, Levy-20D, Plasma-16D): positive across all initializations, which is consistent with the disconnected-graph problem being the main bottleneck in those settings.
- **Low-dimensional benchmarks** (Vehicle-5D, Hartmann-6D, CarCab-7D): mostly neutral. At low dimensionality, the comparison graph provides sufficient curvature relative to the prior, and the correction has little to add.