GUIDE-FBO: GUIDANCE VIA UNCERTAINTY INTERVENTION AND DISTRIBUTIONAL EXCHANGE FOR FEDERATED BAYESIAN OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Federated Bayesian Optimization (FBO) enables distributed agents to collaboratively optimize expensive black-box objectives without sharing raw local observations. However, effective knowledge transfer remains challenging under communication constraints and task heterogeneity. We propose GUIDE-FBO, in which agents exchange compact distributions over the locations of their respective optima inferred from local Gaussian process (GP) posteriors, rather than raw observations, query points, or surrogate parameters. The server merges and reweights these distributional components before returning a subset to each agent. Each agent then constructs a Federated Interventional GP (FI-GP), which preserves the local posterior mean and spatially rescales its covariance for local decision making. For the upper confidence bound (UCB) instantiation, GUIDE-UCB, we prove that any bounded FI-GP uncertainty intervention preserves the leading-order cumulative regret rate of standard GP-UCB. When the transferred distributions place greater support near an optimum than in a suboptimal region, selecting the latter requires greater local posterior uncertainty. Experiments on 12 synthetic benchmarks and three real-world optimization tasks show that GUIDE-FBO remains effective across settings ranging from homogeneous to severely heterogeneous. Ablation results highlight the importance of spatially localized uncertainty intervention, while the communication analysis shows that GUIDE-FBO exchanges only compact distributional messages.

1 INTRODUCTION

Optimizing expensive black-box functions is a ubiquitous challenge across domains ranging from hyperparameter tuning in automated machine learning (AutoML) and drug discovery to advanced manufacturing control. In these settings, the mapping from queries $\mathbf{x}$ to observations y lacks a closed-form expression, and evaluations are costly in terms of time or budget (Shahriari et al., 2016). Bayesian Optimization (BO) is a prominent methodology for such sample-efficient optimization. By employing a probabilistic surrogate model, typically a Gaussian Process (GP) (Seeger, 2004), to approximate the objective function and utilizing an acquisition function to guide the search, BO quantifies uncertainty to strategically balance exploration and exploitation, enabling efficient global optimization within a limited evaluation budget (Wang et al., 2023a).

Recently, BO has been extended to distributed environments characterized by data silos and privacy concerns, termed Federated Bayesian Optimization (FBO) (Dai et al., 2020). We focus on federated multi-task BO. In this context, N agents aim to accelerate the optimization of related but heterogeneous local objective functions $\{f_n(\mathbf{x})\}_{n=1}^N$ through knowledge transfer, without revealing their raw local observations. This setting arises in dose optimization across multiple clinical centers, where evaluations are costly and safety is critical. Differences among institutions and patient populations require personalized decisions (O'Quigley et al., 1990; Yu et al., 2015; Willard et al., 2025). Similar applications appear in federated hyperparameter tuning and robotic controller optimization (Marco et al., 2016; Khodak et al., 2021).

However, transforming this immense potential into practice presents fundamental challenges. An FBO method must keep raw local observations private while addressing two central issues: (I) Task

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.

Heterogeneity: Objective functions can be distinct ($f_n \neq f_m$), meaning blind aggregation risks negative transfer; and (II) Communication Cost: Bandwidth is often limited, necessitating lightweight interaction protocols (Dai et al., 2020; Al Kontar, 2024). Existing methods typically trade off these requirements. Federated Thompson Sampling (FTS) (Dai et al., 2020) and its differentially private distributed exploration variant (DP-FTS-DE) (Dai et al., 2021) exchange random feature representations of local GP surrogates, whose communication cost grows with the approximation dimension. Federated Many-Task BO (FMTBO) (Zhu et al., 2024) communicates lightweight GP hyperparameter information, but such sparse model-level summaries provide limited guidance for local sequential search. Constrained Gaussian Process (CGP) methods (Chen et al., 2025) instead transfer selected high-potential queries. Under task heterogeneity, however, a point that is promising for one task may not be promising for another, making such point-level transfer fragile. These limitations motivate a fundamental question: What information should be exchanged to enable lightweight and effective knowledge transfer across heterogeneous tasks without exposing raw local observations?

To answer this question, we propose GUIDE-FBO, which uses a distributional representation of *where the agents' optima are likely to lie* as the information carrier. Specifically, each agent transfers a distribution over the location of its optimum, as inferred from its GP posterior. The server receives and merges spatially redundant distributional components, reweights the merged set, and sends a subset of the resulting components to each agent. Each agent then converts the received components into a spatially varying scaling function for its local GP posterior covariance, yielding the Federated Interventional GP (FI-GP) for local BO decision making. The underlying philosophy is to direct each agent toward globally supported promising regions, thereby reducing redundant evaluations and improving sample efficiency across the federated system. Meanwhile, communication costs are lowered by transmitting only compact distribution parameters.

Our main contributions are summarized as follows:

- We propose GUIDE-FBO, which exchanges compact distributions over the locations of agents' optima inferred from local GP posteriors. The server merges, reweights, and redistributes these distributions to provide global guidance with lightweight communication.
- We introduce FI-GP, which preserves the local GP posterior mean and spatially rescales its covariance according to the received distributions. It supports standard BO acquisition functions. For GUIDE-UCB, we prove the same leading-order cumulative regret rate as standard GP-UCB under bounded uncertainty intervention and show that informative transferred distributions raise the local uncertainty required to select a suboptimal point.
- GUIDE-FBO achieves the best average rank across all three heterogeneity levels on 12 synthetic benchmarks and remains effective on three real-world tasks. Ablation and communication analyses further support the proposed mechanism and its communication efficiency.

2 Related Work and Problem Setup

2.1 Related Work

We categorize existing federated and collaborative BO methods according to where shared knowledge enters the optimization pipeline.

The first category preserves a local surrogate fitted exclusively to private observations and uses external knowledge to coordinate query selection. In FTS (Dai et al., 2020), agents select query points using their own GPs or GP sample paths transferred from other agents using random Fourier features (Rahimi & Recht, 2007). DP-FTS-DE (Dai et al., 2021) extends this framework with differential privacy and distributed exploration. Beyond sampled function exchange, Yue et al. (2025) combines agents' local query proposals through a consensus matrix whose weights vary over time, gradually shifting from peer collaboration to personalized search. Liu et al. (2026) trains local classifiers on pairwise preferences generated by acquisition functions, aggregates their parameters, and uses the resulting global classifier as an implicit acquisition function.

The second category incorporates shared knowledge into the predictive modeling process used for acquisition optimization. FMTBO (Zhu et al., 2024) shares local GP hyperparameters, estimates task relatedness from predictive rankings, and combines local and aggregated surrogate predictions through a federated ensemble acquisition function. CGP methods (Chen et al., 2025) receive

promising designs from selected collaborators, filter them according to their consistency with local evidence, and use the accepted designs to construct a constrained GP surrogate. Other methods communicate compact symbolic regression models while retaining local GP uncertainty estimates (Wang et al., 2025), or aggregate local GP surrogates through a Wasserstein barycenter to form a central predictive model (Zhan et al., 2025).

GUIDE-FBO differs in how transferred knowledge is integrated into local decision making. Rather than directly coordinating query proposals or conditioning and combining predictive models, GUIDE preserves the local posterior mean and uses the received distributions only to spatially rescale its covariance for the current decision. The resulting FI-GP serves as an auxiliary decision posterior, while the underlying local GP remains fitted exclusively to private observations.

2.2 Federated Multi-Task Bayesian Optimization

We consider a federated system with a central server and N distributed agents, following the standard federated setting (McMahan et al., 2017). Each agent $n \in \{1, \dots, N\}$ owns a black-box objective function $f_n : \mathcal{X} \to \mathbb{R}$ over a shared compact domain $\mathcal{X} \subseteq \mathbb{R}^d$. The local objectives are related but may be heterogeneous. For each agent, let $\mathbf{x}_n^*$ denote an optimum of f_n over $\mathcal{X}$, with

$$\mathbf{x}_n^* \in \arg\max_{\mathbf{x}\in\mathcal{X}} f_n(\mathbf{x}), \qquad n = 1, \dots, N.$$

At the start of round $t \geq 1$, agent n has the private observation set $\mathcal{D}_{n,t-1} = \mathcal{D}_{n,0} \cup \{(\mathbf{x}_{n,i}, y_{n,i})\}_{i=1}^{t-1}$, where $y_{n,i} = f_n(\mathbf{x}_{n,i}) + \epsilon_{n,i}$. The corresponding pre-query local GP posterior is denoted by $p_{n,t-1} := p(f_n \mid \mathcal{D}_{n,t-1}) = \mathcal{GP}(\mu_{n,t-1}, k_{n,t-1})$, with posterior mean $\mu_{n,t-1}$, covariance function $k_{n,t-1}$, and posterior standard deviation $\sigma_{n,t-1}$. After selecting $\mathbf{x}_{n,t}$ and observing $y_{n,t}$, the local dataset is updated as $\mathcal{D}_{n,t} = \mathcal{D}_{n,t-1} \cup \{(\mathbf{x}_{n,t}, y_{n,t})\}$, yielding the updated local posterior $p_{n,t}$. Importantly, raw observations remain local and are never shared with the server or other agents. The goal of federated multi-task BO is to establish a collaborative mechanism that exchanges abstract knowledge through the server and accelerates each agent's search for its optimum $\mathbf{x}_n^*$.

3 Methodology

GUIDE-FBO proceeds in three stages. First, agents upload compact distributions over the locations of their optima. Second, the server aggregates and redistributes these distributional components. Third, each agent uses the received components to construct an FI-GP for local decision making.

At round t, GUIDE-FBO constructs the current guidance from the pre-query local posterior $p_{n,t-1}$. Newly constructed guidance quantities are indexed by t; within a fixed round, the index is omitted when no ambiguity arises.

3.1 Optimum Distribution Extraction

GUIDE-FBO represents transferable knowledge by the posterior distribution over the location of agent n's optimum, denoted by $p(\mathbf{x}_n^* \mid \mathcal{D}_{n,t-1})$. For agent n at round t, since $p(\mathbf{x}_n^* \mid \mathcal{D}_{n,t-1})$ induced by $p_{n,t-1}$ is not available in closed form, we approximate it through posterior sampling. Specifically, we draw M posterior sample paths $\{\widehat{f}_{n,t}^{(m)}\}_{m=1}^M$ from $p_{n,t-1}$ and maximize each sampled function to obtain candidate optimum locations $\widehat{\mathbf{x}}_{n,t}^{(m)} \in \arg\max_{\mathbf{x}\in\mathcal{X}} \widehat{f}_{n,t}^{(m)}(\mathbf{x})$.

The resulting set $\mathcal{C}_{n,t} = \{\widehat{\mathbf{x}}_{n,t}^{(m)}\}_{m=1}^M$ provides a Monte Carlo approximation of where agent n's optimum may lie. We use orthogonal random features to obtain scalable approximate posterior sample paths, with details deferred to Appendix I. To compactly approximate this potentially multimodal distribution over optimum locations, we fit a Dirichlet process Gaussian mixture model (DPGMM) (Rasmussen, 1999) to the sampled optimum locations:

$$q_{n,t}(\mathbf{x}^*) = \sum_{k=1}^{K_{n,t}} \pi_{n,k} \mathcal{N}(\mathbf{x}^* \mid \boldsymbol{\mu}_{n,k}, \boldsymbol{\Sigma}_{n,k}) .$$

Here, $K_{n,t}$ denotes the number of active mixture components at round t. The mixture weight $\pi_{n,k}$ reflects the posterior support assigned to the k-th component, while the mean $\boldsymbol{\mu}_{n,k}$ and covariance

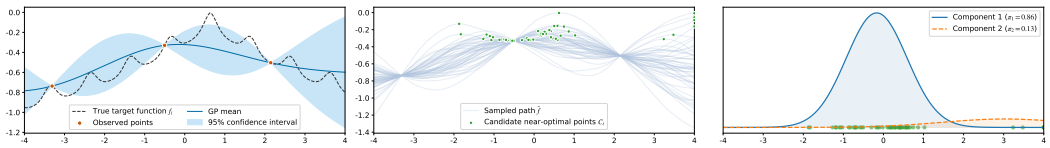

Figure 1: Optimum distribution extraction. A local GP posterior (left) generates candidate optimum locations from posterior sample paths (middle). A DPGMM models their spatial distribution, and its dominant component is selected as the compact distributional message for communication (right).

$\boldsymbol{\Sigma}_{n,k}$ specify its center and spatial spread. This distributional representation converts discrete samples of optimum locations into a compact, transferable summary of promising regions.

To further reduce the communication cost, we use diagonal covariance matrices for the DPGMM components. Under the default uplink budget, GUIDE-FBO uploads only one representative component, selected as the component with the largest mixture weight, $k_n^\dagger \in \arg\max_k \pi_{n,k}$. Figure 1 provides an overview of the complete optimum distribution extraction process.

The mixture weight, however, reflects the posterior support assigned to a component rather than how promising the corresponding region is under the local surrogate. We therefore attach a conservative value score based on the local lower confidence bound (LCB), $\mathrm{LCB}_{n,t} = \mu_{n,t-1}(\boldsymbol{\mu}_{n,k_n^\dagger}) - \kappa\sigma_{n,t-1}(\boldsymbol{\mu}_{n,k_n^\dagger})$, where $\mu_{n,t-1}(\cdot)$ and $\sigma_{n,t-1}(\cdot)$ denote the posterior mean and standard deviation of the local GP before round t, respectively, and $\kappa > 0$ controls the uncertainty penalty. The LCB therefore favors components with high predicted objective values and low local posterior uncertainty.

To account for scale differences across heterogeneous local objectives, we standardize the value estimate as $v_{n,t} = (\mathrm{LCB}_{n,t} - \bar{y}_{n,t-1})/\widehat{\sigma}^{\,y}_{n,t-1}$, where $\bar{y}_{n,t-1}$ and $\widehat{\sigma}^{\,y}_{n,t-1}$ are the empirical mean and standard deviation of the local observations available before round t. Agent n then uploads $\Phi_{n,t} = \left(\pi_{n,k_n^\dagger}, \boldsymbol{\mu}_{n,k_n^\dagger}, \boldsymbol{\Sigma}_{n,k_n^\dagger}, v_{n,t}\right)$ to the server.

3.2 Server-Side Component Merging and Reweighting

At round t, the server receives the uploaded components $\{\Phi_{n,t}\}_{n=1}^N$. Broadcasting all received components would incur a downlink cost that grows with the number of agents. GUIDE-FBO therefore returns only a limited subset to each agent, making it important to avoid spending this budget on spatially redundant components.

The server consequently first merges components that represent nearby regions. Write the received components as $\{(\pi_n, \boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n, v_n)\}_{n=1}^N$. The server measures the distance between two component centers using the root-mean-square (RMS) Euclidean distance $d_{\mathrm{RMS}}(\boldsymbol{\mu}_n, \boldsymbol{\mu}_m) = \sqrt{\|\boldsymbol{\mu}_n - \boldsymbol{\mu}_m\|_2^2/d}$, where $d = \dim(\mathcal{X})$, and performs complete-linkage clustering using these centers. Specifically, the distance between two clusters is $D(\mathcal{I}_a, \mathcal{I}_b) = \max_{n\in\mathcal{I}_a,\, m\in\mathcal{I}_b} d_{\mathrm{RMS}}(\boldsymbol{\mu}_n, \boldsymbol{\mu}_m)$. Starting from singleton clusters, the closest pair is merged whenever $D(\mathcal{I}_a, \mathcal{I}_b) \le \delta_{\mathrm{merge}}$, where $\delta_{\mathrm{merge}} \ge 0$ is the spatial merging threshold, until no eligible pair remains.

Each resulting cluster $\mathcal{I}_\ell$ is represented by a moment-matched Gaussian component:

$$\pi'_\ell = \sum_{n\in\mathcal{I}_\ell} \pi_n, \qquad \boldsymbol{\mu}'_\ell = \frac{\sum_{n\in\mathcal{I}_\ell} \pi_n \boldsymbol{\mu}_n}{\pi'_\ell}, \qquad \boldsymbol{\Sigma}'_\ell = \frac{\sum_{n\in\mathcal{I}_\ell} \pi_n \left[\boldsymbol{\Sigma}_n + (\boldsymbol{\mu}_n - \boldsymbol{\mu}'_\ell)(\boldsymbol{\mu}_n - \boldsymbol{\mu}'_\ell)^\top\right]}{\pi'_\ell}.$$

Its value score is aggregated using the same mixture weighted average, $v'_\ell = (\pi'_\ell)^{-1} \sum_{n\in\mathcal{I}_\ell} \pi_n v_n$. Under the default diagonal covariance implementation, only the diagonal entries of $\boldsymbol{\Sigma}'_\ell$ are retained.

Because π'_ℓ reflects accumulated posterior support but not the estimated quality of the corresponding region, the server further reweights the merged components using their value scores as $\omega_\ell = \pi'_\ell \exp(v'_\ell/\tau_t)/\sum_{j=1}^{L_t} \pi'_j \exp(v'_j/\tau_t)$, where τ_t is the adaptive temperature parameter in the current round. Thus, ω_ℓ balances accumulated posterior support with the conservative value score. Figure 2 summarizes this server-side transformation from the received components to the merged and reweighted global component set.

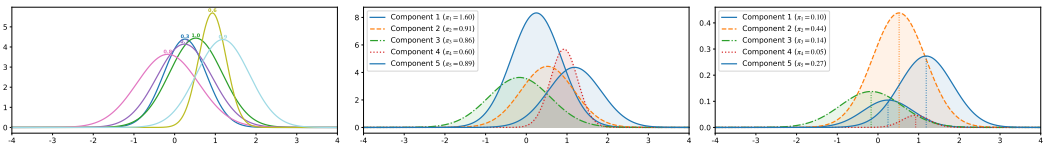

Figure 2: Server-side component processing: received local components (left), moment-matched components after spatial merging (middle), and value-aware reweighted global components (right).

The resulting compact global component set is $\mathcal{B}_t = \{(\omega_\ell, \boldsymbol{\mu}'_\ell, \boldsymbol{\Sigma}'_\ell)\}_{\ell=1}^{L_t}$, where $L_t \leq N$ and ω_ℓ reflects the global importance of each merged component as a promising region. Sending the same subset to every agent could concentrate their parallel searches on the same few regions. To maintain diversity across agents, the server independently samples up to P components without replacement for each agent according to $\{\omega_\ell\}_{\ell=1}^{L_t}$, retaining their global weights. The resulting guidance packet sent to agent n is $\Psi_{n,t} = \{(\omega_{n,p}, \boldsymbol{\mu}_{n,p}, \boldsymbol{\Sigma}_{n,p})\}_{p=1}^{P_t}$, where $P_t = \min\{P, L_t\}$.

3.3 Federated Interventional Gaussian Process

The final stage integrates the received guidance $\Psi_{n,t}$ into agent n's local decision making. Because the exchanged components indicate where promising regions may lie but do not provide objective-value information, FI-GP preserves the local posterior mean and lets federated guidance act only through uncertainty. We refer to this as a mean-preserving uncertainty intervention.

Given $\Psi_{n,t} = \{(\omega_{n,p}, \boldsymbol{\mu}_{n,p}, \boldsymbol{\Sigma}_{n,p})\}_{p=1}^{P_t}$, agent n first constructs a spatial guidance field

$$G_{n,t}(\mathbf{x}) = \sum_{p=1}^{P_t} \omega_{n,p} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_{n,p})^\top \boldsymbol{\Sigma}_{n,p}^{-1}(\mathbf{x} - \boldsymbol{\mu}_{n,p})\right).$$

Each component contributes most strongly around its center, with $\omega_{n,p}$ controlling its influence and $\boldsymbol{\Sigma}_{n,p}$ determining its spatial extent. $G_{n,t}$ is a nonnegative guidance field rather than a probability density. Because each Gaussian-shaped term lies in $(0, 1]$ and the retained global weights satisfy $\sum_{p=1}^{P_t} \omega_{n,p} \leq 1$, we have $0 \leq G_{n,t}(\mathbf{x}) \leq 1$.

We then convert this spatial guidance field into a spatially varying uncertainty scaling function, $S_{n,t}(\mathbf{x}) = 1 + \lambda_t G_{n,t}(\mathbf{x})$, where $\lambda_t = \lambda_{\max}/\sqrt{t}$ controls the intervention strength. It follows that $1 \leq S_{n,t}(\mathbf{x}) \leq 1 + \lambda_t$, so the intervention can only expand local posterior uncertainty. The $1/\sqrt{t}$ decay gives federated guidance greater influence early in optimization and gradually reduces it as local evidence accumulates.

Let $\mu_{n,t-1}(\mathbf{x})$ and $k_{n,t-1}(\mathbf{x}, \mathbf{x}')$ denote the mean and covariance functions of the local GP posterior before round t. FI-GP keeps the local posterior mean unchanged and spatially rescales its covariance to form the round-t decision posterior:

$$\widetilde{\mu}_{n,t}(\mathbf{x}) = \mu_{n,t-1}(\mathbf{x}), \qquad \widetilde{k}_{n,t}(\mathbf{x}, \mathbf{x}') = S_{n,t}(\mathbf{x}) k_{n,t-1}(\mathbf{x}, \mathbf{x}') S_{n,t}(\mathbf{x}').$$

This rescaling preserves positive semidefiniteness and hence defines a valid GP covariance, as shown in Appendix B.3. We denote the resulting FI-GP decision posterior by $\widetilde{p}_{n,t}$. Accordingly, its standard deviation satisfies $\widetilde{\sigma}_{n,t}(\mathbf{x}) = S_{n,t}(\mathbf{x})\sigma_{n,t-1}(\mathbf{x})$. This decision posterior is used only for query selection; the fitted local GP remains unchanged and continues to be updated solely from local observations. Figure 3 illustrates this mean-preserving uncertainty intervention.

Importantly, the absolute effect of federated guidance is further modulated by local posterior uncertainty: $\widetilde{\sigma}_{n,t}(\mathbf{x}) - \sigma_{n,t-1}(\mathbf{x}) = \lambda_t G_{n,t}(\mathbf{x})\sigma_{n,t-1}(\mathbf{x})$. Thus, for the same guidance intensity, the intervention has a smaller absolute effect where the local GP is already more certain.

For acquisition-based rules such as UCB and noisy expected improvement (NEI), agent n selects $\mathbf{x}_{n,t} \in \arg\max_{\mathbf{x}\in\mathcal{X}} \alpha(\mathbf{x}; \widetilde{p}_{n,t})$. For Thompson sampling (TS), the agent applies the same uncertainty scaling to a local posterior sample path and optimizes the resulting path. Thus, the same uncertainty intervention can be combined with different BO decision rules.

The complete GUIDE-FBO procedure is summarized in Algorithm 1 in Appendix A.

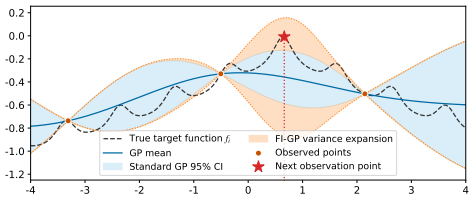

Figure 3: Federated interventional decision making. FI-GP preserves the local posterior mean while expanding uncertainty around regions supported by federated guidance to guide the next query.

4 Theoretical Analysis

We analyze GUIDE-UCB under two guidance regimes. First, under arbitrary bounded guidance, including inaccurate or misleading guidance, we establish whether GUIDE-UCB preserves the regret rate of standard GP-UCB. Second, under informative guidance, we characterize the local posterior uncertainty that a suboptimal point must retain to be selected.

For agent n, let $r_{n,t} = f_n(\mathbf{x}_n^*) - f_n(\mathbf{x}_{n,t})$ denote its instantaneous regret and $R_{n,T} = \sum_{t=1}^{T} r_{n,t}$ its cumulative regret. Using the FI-GP uncertainty scaling defined in Section 3.3, GUIDE-UCB selects

$$\mathbf{x}_{n,t} \in \arg\max_{\mathbf{x}\in\mathcal{X}} \mu_{n,t-1}(\mathbf{x}) + \sqrt{\beta_t}\big(1 + \lambda_t G_{n,t}(\mathbf{x})\big)\sigma_{n,t-1}(\mathbf{x}). \tag{1}$$

Our analysis relies on the standard simultaneous GP confidence event used in GP-UCB and kernelized bandit analyses (Srinivas et al., 2010; Chowdhury & Gopalan, 2017). The guidance field is uniformly bounded as $0 \le G_{n,t}(\mathbf{x}) \le G_{\max}$, with $G_{\max} = 1$ being a valid bound under the normalized weighting used in GUIDE-FBO; see Remark B.1. Formal confidence conditions, auxiliary results, and complete proofs are provided in Appendix B.

Theorem 1 (Cumulative regret under arbitrary bounded guidance)**.** *Under Assumption B.1, suppose that the standard cumulative posterior variance bound in Lemma B.2 holds, $\sum_{t=1}^{T} \sigma_{n,t-1}^2(\mathbf{x}_{n,t}) \le C_\gamma \gamma_{n,T}$, where $\gamma_{n,T}$ is the standard maximum information-gain (log-determinant) complexity term associated with the local kernel of agent n. If $\lambda_t = \lambda_{\max}/\sqrt{t}$ for some fixed $\lambda_{\max} \ge 0$, then, with probability at least $1-\delta$, every agent n satisfies*

$$R_{n,T} \le \sqrt{\beta_T C_\gamma \gamma_{n,T}}\sqrt{4T + 4\lambda_{\max} G_{\max}(2\sqrt{T} - 1) + \lambda_{\max}^2 G_{\max}^2(1 + \log T)}. \tag{2}$$

Consequently, $R_{n,T} = \mathcal{O}\big(\sqrt{T\beta_T\gamma_{n,T}}\big)$.

Theorem 1 shows that GUIDE-UCB preserves the leading-order cumulative regret rate of standard GP-UCB under arbitrary bounded guidance. In Eq. 2, the guidance-dependent terms grow only as $\sqrt{T}$ and $\log T$, compared with the leading T term inside the square root. Thus, bounded guidance affects the regret bound only through lower-order terms. A fixed intervention strength $\lambda_t \equiv \lambda$ also preserves the same leading-order regret rate up to a constant factor; see Appendix B.6.

Proposition 1 (Selection threshold under informative guidance)**.** *Under Assumption B.1, fix an agent n and round t, and suppose that $\lambda_t \ge 0$. If GUIDE-UCB selects a suboptimal point $\mathbf{x}_{n,t}$ with $\Delta_n(\mathbf{x}_{n,t}) := f_n(\mathbf{x}_n^*) - f_n(\mathbf{x}_{n,t}) > 0$, then*

$$\sigma_{n,t-1}(\mathbf{x}_{n,t}) \ge \frac{\Delta_n(\mathbf{x}_{n,t}) + \lambda_t\sqrt{\beta_t} G_{n,t}(\mathbf{x}_n^*)\sigma_{n,t-1}(\mathbf{x}_n^*)}{\sqrt{\beta_t}\big(2 + \lambda_t G_{n,t}(\mathbf{x}_{n,t})\big)}. \tag{3}$$

If, in addition, $G_{n,t}(\mathbf{x}_n^*) \ge c$ and $G_{n,t}(\mathbf{x}_{n,t}) \le \varepsilon_G$ for some $c > 0$ and $\varepsilon_G \ge 0$, then under $\lambda_t = \lambda_{\max}/\sqrt{t}$ with $\lambda_{\max} \ge 0$,

$$\sigma_{n,t-1}(\mathbf{x}_{n,t}) \ge \frac{\Delta_n(\mathbf{x}_{n,t}) + \frac{\lambda_{\max} c}{\sqrt{t}}\sqrt{\beta_t}\sigma_{n,t-1}(\mathbf{x}_n^*)}{\sqrt{\beta_t}\left(2 + \frac{\lambda_{\max}\varepsilon_G}{\sqrt{t}}\right)}.$$

Proposition 1 gives a necessary local uncertainty condition for selecting a suboptimal point under informative guidance. When the guidance is large near the agent's optimum and small at a suboptimal

candidate, that candidate can be selected only if it retains sufficiently large local posterior uncertainty. In particular, when ε_G is negligible, the lower bound on the required uncertainty exceeds the standard GP-UCB counterpart $\Delta_n(\mathbf{x}_{n,t})/(2\sqrt{\beta_t})$ by approximately $\lambda_{\max} c\, \sigma_{n,t-1}(\mathbf{x}_n^*)/(2\sqrt{t})$. Thus, when guidance favors the optimum over a suboptimal candidate, the latter can be selected only if it remains sufficiently uncertain under the local GP.

5 Experiments

5.1 Experimental Setup

Baselines. We compare GUIDE-FBO against nine competitive methods from three categories. The first category contains non-collaborative BO methods, including TS (Thompson, 1933), UCB (Srinivas et al., 2010), and NEI (Letham et al., 2019; Balandat et al., 2020). The second category contains classical FBO methods, including FTS (Dai et al., 2020) and its distributed exploration variant FTS-DE (Dai et al., 2021). The third category contains recent federated and collaborative BO methods, including FMTBO (Zhu et al., 2024) and the CGP family (CGP-TS, CGP-UCB, and CGP-NEI) (Chen et al., 2025). We instantiate GUIDE-FBO with the same three local acquisition rules, yielding GUIDE-TS, GUIDE-UCB, and GUIDE-NEI. These variants use the same distributional exchange and uncertainty intervention mechanism and differ only in the local acquisition function.

Benchmarks and heterogeneity. We evaluate all methods on 12 synthetic benchmark functions and three real-world tasks. Landmine Detection (Xue et al., 2007) uses 29 landmine fields as agents that tune support vector machines with radial basis function kernels; Activity Recognition (Anguita et al., 2013) uses 30 subjects as agents that tune logistic-regression classifiers; and FedHPO-Bench (FedHPO) (Wang et al., 2023b) uses Cora, CiteSeer, and PubMed as three graph convolutional network hyperparameter-optimization tasks evaluated with the official surrogates. For each synthetic benchmark and experimental run, the local objectives are generated from a common base function using independently sampled random shifts and rotations for different agents. The same realized objectives are used by all compared methods within each run. Level 1 uses no transformation; Level 2 uses $(\delta_{\text{shift}}, \delta_{\text{rot}}) = (0.05, 0.1)$; and Level 3 uses $(0.3, 1.0)$, corresponding to homogeneous, mild, and severe heterogeneity, respectively. Complete task construction and heterogeneity details are provided in Appendix C.3.

Protocol and metrics. All experiments use 30 initial evaluations per agent and 50 BO rounds and are repeated over 10 independent runs. Synthetic experiments use $N = 16$ agents with Gaussian observation noise of standard deviation 0.1, while the real-world tasks use their natural numbers of agents. For synthetic benchmarks, we report simple regret averaged across agents. For Landmine Detection, we report the area under the receiver operating characteristic curve (AUC), while Activity Recognition and FedHPO use validation accuracy. Higher values are better for all real-world metrics. Complete benchmark definitions and implementation details are provided in Appendix C.

5.2 Main Results

Synthetic benchmarks. Table 1 shows that GUIDE achieves strong aggregate performance across all three heterogeneity levels. GUIDE-UCB obtains the best average rank over the 12 synthetic benchmarks at Levels 1–3, with ranks of 1.33, 1.58, and 2.58, respectively, while GUIDE-NEI ranks second with 2.42, 2.17, and 3.42. Matched acquisition comparisons further isolate the contribution of GUIDE itself: GUIDE-UCB improves upon UCB on $12/12$, $11/12$, and $8/12$ benchmarks from Levels 1–3; GUIDE-NEI improves upon NEI on $12/12$, $11/12$, and $9/12$; and GUIDE-TS improves upon TS on $12/12$, $12/12$, and $8/12$. These results show that the benefit of GUIDE extends across different acquisition rules rather than being specific to one of them.

At Levels 1 and 2, where the empirical mean pairwise Spearman correlations of objective rankings are 1.000 and 0.557, respectively, GUIDE-UCB and GUIDE-NEI frequently reach lower-regret regions earlier and maintain favorable convergence trajectories as more local observations are collected. This pattern reflects the mechanism of GUIDE: FI-GP increases local posterior uncertainty around regions supported by the transferred distributions, allowing them to be explored earlier when they remain plausible under the local posterior.

At Level 3, the mean pairwise Spearman correlation of objective rankings decreases to 0.081, indicating that useful information is substantially less transferable across agents. GUIDE-UCB and GUIDE-NEI nevertheless retain the first and second best average ranks, although their advantages become smaller and more task dependent. Two design features limit the influence of mismatched guidance. First, the intervention strength decays over time, so federated guidance gradually gives way to local evidence. Second, the uncertainty increment is proportional to local posterior uncertainty, so guidance has little effect where the local GP is already confident. Negative transfer can nevertheless occur on individual tasks under severe heterogeneity. Detailed results and convergence analyses are provided in Appendix D.

Table 1: Average rank across 12 synthetic benchmarks under three heterogeneity levels, computed from the final average simple regret over 10 independent runs. Lower is better; best and second-best results are shown in bold and underlined.

Heterogeneity	TS	UCB	NEI	FTS	FTS-DE	FMTBO	CGP-TS	CGP-UCB	CGP-NEI	GUIDE-TS	GUIDE-UCB	GUIDE-NEI
Level 1: $(0, 0)$	11.33	7.00	5.50	10.42	9.83	5.75	10.17	4.25	2.75	7.25	**1.33**	<u>2.42</u>
Level 2: $(0.05, 0.1)$	10.67	6.17	5.33	10.25	10.75	5.58	10.25	4.67	3.08	7.50	**1.58**	<u>2.17</u>
Level 3: $(0.3, 1.0)$	9.17	4.83	5.42	11.08	10.33	3.67	8.33	4.25	6.75	8.17	**2.58**	<u>3.42</u>

Real-world tasks. Table 2 shows that a GUIDE variant achieves the highest final mean performance on each real-world task, although the differences among the leading methods are small on Activity Recognition and FedHPO. GUIDE-UCB obtains the highest AUC on Landmine Detection, while GUIDE-NEI obtains the highest validation accuracy on Activity Recognition and FedHPO and the best average rank of 1.67. Moreover, GUIDE-UCB has higher final mean performance than UCB, and GUIDE-NEI than NEI, on all three tasks. By contrast, GUIDE-TS improves over TS only on Landmine Detection and has an average rank of 9.67 across the three tasks. Complete convergence trajectories are provided in Appendix D.4.

Table 2: Results on three real-world federated optimization benchmarks, averaged over 10 independent runs. Higher task performance and lower average rank are better; best and second-best results are shown in bold and underlined.

Task	TS	UCB	NEI	FTS	FTS-DE	FMTBO	CGP-TS	CGP-UCB	CGP-NEI	GUIDE-TS	GUIDE-UCB	GUIDE-NEI
Landmine Detection (AUC)	0.803987	0.805593	0.805508	0.804052	0.801349	0.805970	0.805247	<u>0.806501</u>	0.805218	0.805001	**0.806604**	0.806452
Activity Recognition (Acc.)	0.987787	0.987821	0.987804	0.987800	0.987739	0.987891	0.987701	0.987959	<u>0.987990</u>	0.987754	0.987938	**0.987994**
FedHPO (Acc.)	0.850968	0.850973	0.851145	0.850439	0.851098	0.851049	0.850766	<u>0.851262</u>	0.851068	0.850853	0.851178	**0.851271**
Average rank	9.67	6.33	5.67	10.00	9.33	5.33	10.00	<u>2.33</u>	5.33	9.67	2.67	**1.67**

5.3 Ablation and Communication Efficiency

Ablation study. Table 3 examines the contribution of the main components of GUIDE-UCB. Under Level 2 heterogeneity, Full GUIDE-UCB achieves the best average rank of 1.50 and remains in the top two on all four representative benchmarks. The largest degradation occurs when the spatially varying uncertainty intervention is replaced by uniform uncertainty scaling: its average rank increases to 6.25, close to Vanilla UCB at 6.75. The absolute results show the same pattern. On Ackley, Full GUIDE-UCB, Uniform uncertainty scaling, and Vanilla UCB obtain 8.4785, 19.1366, and 19.4350, respectively; on Rastrigin, the corresponding values are 58.3082, 80.3952, and 80.4771. These results indicate that, among the tested Level 2 ablations, the spatial localization of the uncertainty intervention is the strongest factor distinguishing GUIDE-UCB from a uniform increase in posterior uncertainty.

The effects of the remaining modules are smaller and more task dependent. Multimodal belief modeling, value-aware reweighting, component merging, and agent-specific sampling contribute to the aggregate performance of the complete method, although individual ablations can outperform Full GUIDE-UCB on particular functions. Appendix E further shows that the relative importance of these components changes under stronger heterogeneity. At Level 3, Full GUIDE-UCB retains the best aggregate rank of 2.50, whereas removing value-aware reweighting gives the worst rank of 5.25. This suggests that value-aware server reweighting becomes more important when cross-agent agreement is weak.

Table 3: Ablation study of GUIDE-UCB on four synthetic benchmarks under Level 2 heterogeneity. Final average simple regret over 10 runs is reported; lower is better. The best and second-best results in each row are shown in bold and underlined, respectively.

Benchmark	Vanilla UCB	Single-Gaussian belief	w/o value-aware reweighting	w/o component merging	w/o agent-specific sampling	Uniform uncertainty scaling	Full GUIDE-UCB
Ackley	19.4350	8.6510	**8.4753**	8.5349	8.4855	19.1366	<u>8.4785</u>
Rastrigin	80.4771	60.2355	<u>59.1892</u>	59.9796	60.0268	80.3952	**58.3082**
Zakharov	49.8750	25.5947	22.9307	25.3701	**21.5431**	52.0749	<u>22.1717</u>
Michalewicz	6.1248	5.9494	5.9814	5.9650	<u>5.8771</u>	6.1078	**5.8629**
Average rank	6.75	4.50	2.75	3.75	<u>2.50</u>	6.25	**1.50**

Communication efficiency. Table 4 shows that GUIDE-FBO obtains its optimization gains with a compact communication message. Under the default synthetic setting with $d = 10$ and $P = 5$, GUIDE-FBO communicates at most 127 scalars per agent per BO round, compared with 212 for CGP, 1000 for FTS, and 4000 for FTS-DE. FMTBO has the lowest communication cost among the collaborative baselines, requiring only 14 scalars per round. This lower communication volume, however, is accompanied by weaker aggregate optimization performance than GUIDE-UCB on the synthetic benchmarks, while both GUIDE-UCB and GUIDE-NEI achieve higher final mean performance than FMTBO on all three real-world tasks. Overall, GUIDE provides a favorable trade-off between preserving spatial information about promising regions and limiting communication cost.

Table 4: Communication cost per agent per round for $d = 10$. GUIDE-FBO uses diagonal covariance and downloads at most $P = 5$ components per round.

	FTS	FTS-DE	FMTBO	CGP	GUIDE
Uplink	500	2000	12	12	22
Downlink	500	2000	2	200	≤ 105
Total	1000	4000	14	212	≤ 127

5.4 Additional Analyses

The additional experiments further examine the sensitivity and computational efficiency of GUIDE. The Level 3 sensitivity study in Appendix F shows that, on the tested benchmarks, GUIDE remains stable across moderate choices of M, P, and δ_{merge}, while overly strong uncertainty intervention degrades performance. These results suggest that the default configuration is not narrowly tuned to a single parameter setting. The runtime study in Appendix H shows that GUIDE introduces additional computation relative to independent BO, but remains more efficient than the corresponding CGP variants on most Level 2 synthetic benchmarks. In applications with expensive objective evaluations, this optimizer-side overhead may constitute a smaller fraction of the total optimization cost.

6 Conclusion

GUIDE-FBO approaches federated Bayesian optimization from the perspective of what should be communicated across heterogeneous agents. By representing local search information as distributions over optimum locations, it separates global knowledge exchange from local surrogate fitting and uses transferred information only to spatially rescale posterior uncertainty during local decision making. This intervention can be combined with UCB, NEI, and TS without redesigning their local decision rules. Our analysis of GUIDE-UCB establishes the same leading-order regret rate as GP-UCB under bounded federated guidance and shows that, when guidance is sufficiently stronger near the optimum than at a suboptimal point, the latter can be selected only if it remains sufficiently uncertain under the local GP. Experiments further demonstrate the value of this mechanism across heterogeneous tasks. Although GUIDE-FBO avoids sharing raw local observations, the current framework does not provide a formal privacy guarantee for the exchanged distributional information. Future work could incorporate differential privacy into local distribution release or server aggregation while studying how privacy noise interacts with federated guidance.

AI USE STATEMENT

In this work, we used generative AI tools for the following tasks with required disclosure:

- Provide critical ingredients for proving mathematical claims
- Design or provide feedback on research methodology or experiments
- Implement methods
- Assist with translation

We have not used generative AI tools for the following tasks with required disclosure:

- Generate synthetic data sets
- Help develop theoretical models or conceptual frameworks
- Formulate mathematical claims
- Assist in the writing of proofs
- Propose or refine hypotheses
- Clean and reformat dataset
- Interpret results

The following task with required disclosure is not applicable to this work:

- Support qualitative and thematic data analysis

Additionally, we used generative AI tools for the following tasks with recommended disclosure:

- Create or edit software code
- Sourcing/searching for information
- Edit a research paper to improve readability
- Propose a title or keywords for a research paper

We have reviewed all AI-assisted work. Mathematical claims and proofs were checked by the authors, experimental interpretations were verified against the reported results, AI-assisted code was reviewed and tested.

We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

Detailed experimental settings, benchmark construction, data processing, and GUIDE-FBO hyperparameters are provided in Appendix C. Complete source code and experiment scripts are included in the anonymized supplementary material. Full assumptions and proofs for the theoretical results are given in Appendix B.

REFERENCES

Raed Al Kontar. Collaborative and federated black-box optimization: A Bayesian optimization perspective. In *2024 IEEE International Conference on Big Data (BigData)*, pp. 7854–7859. IEEE, 2024.

Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra, and Jorge Luis Reyes-Ortiz. A public domain dataset for human activity recognition using smartphones. In *Proceedings of the 21st European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, pp. 437–442, 2013.

Maximilian Balandat, Brian Karrer, Daniel Jiang, Samuel Daulton, Ben Letham, Andrew G Wilson, and Eytan Bakshy. BoTorch: A framework for efficient monte-carlo Bayesian optimization. In *Advances in Neural Information Processing Systems*, volume 33, pp. 21524–21538, 2020.

Qiyuan Chen, Liangkui Jiang, Hantang Qin, and Raed Al Kontar. Multi-agent collaborative Bayesian optimization via constrained Gaussian processes. *Technometrics*, 67(1):32–45, 2025.

Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 844–853. PMLR, 2017.

Zhongxiang Dai, Bryan Kian Hsiang Low, and Patrick Jaillet. Federated Bayesian optimization via Thompson sampling. In *Advances in Neural Information Processing Systems*, volume 33, pp. 9687–9699, 2020.

Zhongxiang Dai, Bryan Kian Hsiang Low, and Patrick Jaillet. Differentially private federated Bayesian optimization with distributed exploration. In *Advances in Neural Information Processing Systems*, volume 34, pp. 9125–9139, 2021.

Nikolaus Hansen, Steffen Finck, Raymond Ros, and Anne Auger. Real-Parameter Black-Box Optimization Benchmarking 2009: Noiseless Functions Definitions. Technical Report RR-6829, INRIA, 2009.

Nikolaus Hansen, Anne Auger, Raymond Ros, Olaf Mersmann, Tea Tušar, and Dimo Brockhoff. COCO: A platform for comparing continuous optimizers in a black-box setting. *Optimization Methods and Software*, 36(1):114–144, 2021.

Mikhail Khodak, Renbo Tu, Tian Li, Liam Li, Maria-Florina F Balcan, Virginia Smith, and Ameet Talwalkar. Federated hyperparameter tuning: Challenges, baselines, and connections to weight-sharing. In *Advances in Neural Information Processing Systems*, volume 34, pp. 19184–19197, 2021.

Benjamin Letham, Brian Karrer, Guilherme Ottoni, and Eytan Bakshy. Constrained Bayesian Optimization with Noisy Experiments. *Bayesian Analysis*, 14(2):495–519, 2019.

Qiqi Liu, Yaochu Jin, and Guodong Chen. Optimization of an implicit acquisition function for federated Bayesian many-task optimization. *IEEE Trans. Evol. Comput.*, 30(3):1068–1082, 2026.

Alonso Marco, Philipp Hennig, Jeannette Bohg, Stefan Schaal, and Sebastian Trimpe. Automatic LQR tuning based on Gaussian process global optimization. In *2016 IEEE international conference on robotics and automation (ICRA)*, pp. 270–277. IEEE, 2016.

Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pp. 1273–1282, 2017.

John O'Quigley, Margaret Pepe, and Lloyd Fisher. Continual reassessment method: A practical design for phase 1 clinical trials in cancer. *Biometrics*, 46(1):33–48, 1990.

Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, volume 20, pp. 1177–1184, 2007.

Carl Edward Rasmussen. The infinite Gaussian mixture model. In *Advances in Neural Information Processing Systems*, volume 12, pp. 554–560. MIT Press, 1999.

Matthias Seeger. Gaussian processes for machine learning. *International Journal of Neural Systems*, 14(2):69–106, 2004.

Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1): 148–175, 2016.

Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias W. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the 27th International Conference on Machine Learning*, pp. 1015–1022, 2010.

William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.

Xilu Wang, Yaochu Jin, Sebastian Schmitt, and Markus Olhofer. Recent advances in Bayesian optimization. *ACM Computing Surveys*, 55(13s):1–36, 2023a.

Xilu Wang, Kaifeng Yang, Peng Liao, Mengxuan Zhang, and Yaochu Jin. Efficient federated Bayesian optimization with symbolic regression model. In *2025 IEEE Congress on Evolutionary Computation (CEC)*, pp. 1–9. IEEE, 2025.

Zhen Wang, Weirui Kuang, Ce Zhang, Bolin Ding, and Yaliang Li. FedHPO-bench: A benchmark suite for federated hyperparameter optimization. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 35908–35948, 2023b.

James Willard, Shirin Golchi, Erica EM Moodie, Bruno Boulanger, and Bradley P Carlin. Bayesian optimization for personalized dose-finding trials with combination therapies. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 74(2):373–390, 2025.

Ya Xue, Xuejun Liao, Lawrence Carin, and Balaji Krishnapuram. Multi-task learning for classification with Dirichlet process priors. *Journal of Machine Learning Research*, 8(2):35–63, 2007.

Felix X. Yu, Ananda Theertha Suresh, Krzysztof Marcin Choromanski, Daniel N. Holtmann-Rice, and Sanjiv Kumar. Orthogonal random features. In *Advances in Neural Information Processing Systems*, volume 29, pp. 1975–1983, 2016.

Shipeng Yu, Faisal Farooq, Alexander van Esbroeck, Glenn Fung, Vikram Anand, and Balaji Krishnapuram. Predicting readmission risk with institution-specific prediction models. *Artificial Intelligence in Medicine*, 65(2):89–96, 2015.

Xubo Yue, Yang Liu, Albert S Berahas, Blake N Johnson, and Raed Al Kontar. Collaborative and distributed Bayesian optimization via consensus. *IEEE Transactions on Automation Science and Engineering*, 22:11343–11355, 2025.

Donglin Zhan, Haoting Zhang, Rhonda Righter, Zeyu Zheng, and James Anderson. Collaborative Bayesian optimization via Wasserstein barycenters. In *2025 IEEE 64th Conference on Decision and Control (CDC)*, pp. 6284–6291. IEEE, 2025.

Hangyu Zhu, Xilu Wang, and Yaochu Jin. Federated many-task Bayesian optimization. *IEEE Trans. Evol. Comput.*, 28(4):980–993, 2024.

APPENDIX CONTENTS

A. Complete GUIDE-FBO Pseudocode

B. Theoretical Assumptions, Lemmas, and Proofs

C. Benchmarks, Heterogeneity, Metrics, and Experimental Settings

D. Complete Experimental Results and Convergence Trajectories

E. Ablation Study Details

F. Sensitivity Analysis

G. Communication Cost Analysis

H. Computational Efficiency Analysis

I. Orthogonal Random Features for Posterior-Optimum Sampling

A COMPLETE GUIDE-FBO PSEUDOCODE

The complete optimization procedure is given in Algorithm 1. The three stages correspond directly to optimum distribution extraction, server-side component merging and reweighting, and FI-GP-based local decision making.

Algorithm 1 GUIDE-FBO

Require: Agents $\{1, \dots, N\}$, initial local datasets $\{\mathcal{D}_{n,0}\}_{n=1}^{N}$, total rounds T, posterior-optimum samples M, downlink packet size P, merge threshold δ_{merge}, guidance strength $\lambda_{\max}$

Ensure: Best observed local incumbents $\{\mathbf{x}_n^{\mathrm{best}}\}_{n=1}^{N}$

for each agent n in parallel **do**
  Fit the initial local GP posterior $p_{n,0} := p(f_n \mid \mathcal{D}_{n,0})$
end for
for $t = 1, \dots, T$ **do**
  Optimum distribution extraction
  for each agent n in parallel **do**
    Draw M posterior sample paths from $p_{n,t-1}$ and collect candidate optima $\mathcal{C}_{n,t}$
    Fit a DPGMM $q_{n,t}(\mathbf{x}^*)$ to $\mathcal{C}_{n,t}$
    Select $k_n^\dagger \in \arg\max_k \pi_{n,k}$ and compute the standardized conservative value score $v_{n,t}$
    Upload $\Phi_{n,t} = (\pi_{n,k_n^\dagger}, \boldsymbol{\mu}_{n,k_n^\dagger}, \boldsymbol{\Sigma}_{n,k_n^\dagger}, v_{n,t})$
  end for
  Server-side component merging and reweighting
  Cluster uploaded components using complete linkage with threshold δ_{merge}
  Moment-match each cluster $\mathcal{I}_\ell$ to obtain $(\pi'_\ell, \boldsymbol{\mu}'_\ell, \boldsymbol{\Sigma}'_\ell, v'_\ell)_{\ell=1}^{L_t}$
  Compute $\omega_\ell \propto \pi'_\ell \exp(v'_\ell/\tau_t)$ and form $\mathcal{B}_t = \{(\omega_\ell, \boldsymbol{\mu}'_\ell, \boldsymbol{\Sigma}'_\ell)\}_{\ell=1}^{L_t}$
  for each agent n **do**
    Set $P_t = \min\{P, L_t\}$
    Sample P_t components without replacement from $\mathcal{B}_t$ according to $\{\omega_\ell\}_{\ell=1}^{L_t}$
    Send $\Psi_{n,t} = \{(\omega_{n,p}, \boldsymbol{\mu}_{n,p}, \boldsymbol{\Sigma}_{n,p})\}_{p=1}^{P_t}$
  end for
  Federated interventional decision making
  for each agent n in parallel **do**
    Construct $G_{n,t}(\mathbf{x})$ and $S_{n,t}(\mathbf{x}) = 1 + \lambda_t G_{n,t}(\mathbf{x})$, where $\lambda_t = \lambda_{\max}/\sqrt{t}$
    Construct $\widetilde{p}_{n,t}$ by preserving the mean of $p_{n,t-1}$ and spatially scaling its covariance
    Select $\mathbf{x}_{n,t}$ using the chosen local BO decision rule under $\widetilde{p}_{n,t}$
    Evaluate $y_{n,t} = f_n(\mathbf{x}_{n,t}) + \epsilon_{n,t}$
    Update $\mathcal{D}_{n,t} = \mathcal{D}_{n,t-1} \cup \{(\mathbf{x}_{n,t}, y_{n,t})\}$ and refit $p_{n,t}$
  end for
end for
return $\mathbf{x}_n^{\mathrm{best}} \in \arg\max_{(\mathbf{x},y)\in\mathcal{D}_{n,T}} y$ for each agent n

B THEORETICAL ASSUMPTIONS, LEMMAS, AND PROOFS

This appendix states the confidence conditions and structural properties used in Section 4, establishes the auxiliary results required by the analysis, and provides the complete proofs. Throughout, round t starts from the local posterior conditioned on $\mathcal{D}_{n,t-1}$, while all guidance and decision quantities constructed during that round carry index t. For theoretical clarity, we additionally restore the round index on all received guidance-component parameters in this appendix.

Assumption B.1 (Local GP confidence event)**.** Let $\{\beta_t\}_{t\geq 1}$ be a positive and nondecreasing sequence. With probability at least $1-\delta$, the following event holds simultaneously for all agents $n \in [N]$, rounds $t \geq 1$, and points $\mathbf{x} \in \mathcal{X}$:

$$|f_n(\mathbf{x}) - \mu_{n,t-1}(\mathbf{x})| \leq \sqrt{\beta_t}\sigma_{n,t-1}(\mathbf{x}).$$

Assumption B.1 is the standard simultaneous confidence event used in GP-UCB and kernelized bandit analyses (Srinivas et al., 2010; Chowdhury & Gopalan, 2017). Sufficient RKHS and noise conditions are provided in Section B.1.

Remark B.1 (Nonnegativity and boundedness of the GUIDE field). For agent n at round t, the GUIDE field is

$$G_{n,t}(\mathbf{x}) = \sum_{p=1}^{P_t} \omega_{n,t,p} \exp\left(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu}_{n,t,p})^\top \boldsymbol{\Sigma}_{n,t,p}^{-1}(\mathbf{x}-\boldsymbol{\mu}_{n,t,p})\right),$$

where $\omega_{n,t,p} \geq 0$ and $\boldsymbol{\Sigma}_{n,t,p} \succ 0$. Each Gaussian guidance factor lies in $(0,1]$, and therefore

$$0 \leq G_{n,t}(\mathbf{x}) \leq \sum_{p=1}^{P_t} \omega_{n,t,p}.$$

Under the normalized global reweighting and retained-weight downlink protocol,

$$\sum_{p=1}^{P_t} \omega_{n,t,p} \leq 1$$

for every agent n and round t. Hence,

$$G_{\max} := \sup_{n\in[N],\, t\geq 1,\, \mathbf{x}\in\mathcal{X}} G_{n,t}(\mathbf{x}) \leq 1.$$

Lemma B.1 (Instantaneous regret decomposition)**.** *Under Assumption B.1, for any nonnegative intervention strength λ_t, the GUIDE-UCB rule in Eq. 1 satisfies, for every agent n and round t,*

$$r_{n,t} \leq \sqrt{\beta_t}\Big[2\sigma_{n,t-1}(\mathbf{x}_{n,t}) + \lambda_t G_{n,t}(\mathbf{x}_{n,t})\sigma_{n,t-1}(\mathbf{x}_{n,t}) - \lambda_t G_{n,t}(\mathbf{x}_n^*)\sigma_{n,t-1}(\mathbf{x}_n^*)\Big].$$

Lemma B.1 is the one-step inequality underlying both theoretical results. For the cumulative regret bound, the final nonpositive term is discarded. For the informative-guidance result, it is retained to characterize how guidance around $\mathbf{x}_n^*$ changes the selection threshold of suboptimal points.

B.1 SUFFICIENT CONDITIONS FOR THE LOCAL CONFIDENCE EVENT

The main analysis is stated directly in terms of Assumption B.1. A standard set of sufficient conditions is given below.

Assumption B.2 (Compact domain and bounded kernels)**.** The decision domain $\mathcal{X}$ is compact. For every agent n, the kernel k_n is bounded on the diagonal:

$$k_n(\mathbf{x},\mathbf{x}) \leq 1, \qquad \forall \mathbf{x} \in \mathcal{X}.$$

Assumption B.3 (RKHS-bounded local objectives)**.** For every agent n, the objective f_n belongs to the RKHS $\mathcal{H}_{k_n}$ induced by k_n and satisfies

$$\|f_n\|_{\mathcal{H}_{k_n}} \leq B_n.$$

Let $B = \max_{n\in[N]} B_n$.

Assumption B.4 (Sub-Gaussian observation noise)**.** At round t, agent n observes

$$y_{n,t} = f_n(\mathbf{x}_{n,t}) + \epsilon_{n,t},$$

where $\epsilon_{n,t}$ is conditionally R_ϵ-sub-Gaussian with respect to the filtration generated by the observation history.

Under Assumptions B.2–B.4, standard kernelized-bandit concentration results imply Assumption B.1 for an appropriate nondecreasing confidence sequence $\{\beta_t\}_{t\geq 1}$ (Srinivas et al., 2010; Chowdhury & Gopalan, 2017). A union bound over the N agents yields simultaneous validity across the federation.

As is standard in kernelized-bandit analyses, the theoretical local posterior is formed using a fixed kernel k_n and a fixed positive regularization parameter $\sigma^2_{\text{gp}} > 0$. This parameter determines the posterior variance and the log-determinant complexity term used below; it does not impose a Gaussian distribution on the actual observation noise, which remains conditionally R_ϵ-sub-Gaussian as specified in Assumption B.4. In the empirical implementation, kernel hyperparameters are estimated from the local observations.

B.2 Maximum Information Gain

The optimization procedure starts from a nonempty initial local dataset $\mathcal{D}_{n,0}$. To keep the theoretical notation consistent with the round convention in Section 2.2, we condition the sequential analysis on this initial dataset.

Let $\mathcal{X}_{n,0} = \{\mathbf{x} : (\mathbf{x}, y) \in \mathcal{D}_{n,0}\}$ denote the initial design locations of agent n. For the fixed theoretical kernel k_n and GP regularization parameter $\sigma^2_{\text{gp}} > 0$, define the covariance function after conditioning on the initial dataset as

$$\begin{aligned} k_n^{(0)}(\mathbf{x}, \mathbf{x}') = k_n(\mathbf{x}, \mathbf{x}') \\ - k_n(\mathbf{x}, \mathcal{X}_{n,0}) \left(\mathbf{K}_{n,0} + \sigma^2_{\text{gp}}\mathbf{I}\right)^{-1} k_n(\mathcal{X}_{n,0}, \mathbf{x}'), \end{aligned}$$

where

$$\mathbf{K}_{n,0} = \left[k_n(\mathbf{x}, \mathbf{x}')\right]_{\mathbf{x},\mathbf{x}' \in \mathcal{X}_{n,0}}.$$

Thus, $k_n^{(0)}$ is precisely the covariance of the local GP posterior available before the first BO round. Subsequent posterior variances $\sigma^2_{n,t-1}(\mathbf{x})$ are obtained by further conditioning this initial posterior on the BO observations collected in rounds $1, \ldots, t-1$.

For a finite query set $A = \{\mathbf{x}_1, \ldots, \mathbf{x}_T\} \subset \mathcal{X}$, define

$$\mathbf{K}^{(0)}_{n,A} = \left[k_n^{(0)}(\mathbf{x}, \mathbf{x}')\right]_{\mathbf{x},\mathbf{x}' \in A}.$$

The maximum information gain conditional on the initial dataset is

$$\gamma^{(0)}_{n,T} = \frac{1}{2} \max_{A \subset \mathcal{X}:\, |A| = T} \log\det\left(\mathbf{I} + \sigma^{-2}_{\text{gp}} \mathbf{K}^{(0)}_{n,A}\right).$$

For comparison, let

$$\gamma_{n,T} = \frac{1}{2} \max_{A \subset \mathcal{X}:\, |A| = T} \log\det\left(\mathbf{I} + \sigma^{-2}_{\text{gp}} \mathbf{K}_{n,A}\right),$$

where

$$\mathbf{K}_{n,A} = \left[k_n(\mathbf{x}, \mathbf{x}')\right]_{\mathbf{x},\mathbf{x}' \in A}.$$

This is the standard maximum information-gain, or equivalently log-determinant complexity, quantity used in GP-UCB and kernelized-bandit analyses (Srinivas et al., 2010; Chowdhury & Gopalan, 2017).

Conditioning on the initial observations cannot increase covariance. Therefore, for every finite A,

$$\mathbf{K}^{(0)}_{n,A} \preceq \mathbf{K}_{n,A},$$

which implies

$$\gamma_{n,T}^{(0)} \le \gamma_{n,T}. \tag{4}$$

Consequently, the standard information-gain complexity $\gamma_{n,T}$ remains a valid upper bound after accounting for the initial local dataset.

Under a Gaussian observation model, the log-determinant quantities above coincide with the corresponding mutual information between noisy observations and latent function values. The analysis here only requires their log-determinant forms and therefore does not require Gaussian observation noise.

Lemma B.2 (Cumulative posterior variance)**.** *Under Assumption B.2, for any adaptively selected sequence $\{\mathbf{x}_{n,t}\}_{t=1}^{T}$ generated after conditioning on $\mathcal{D}_{n,0}$,*

$$\sum_{t=1}^{T} \sigma_{n,t-1}^2(\mathbf{x}_{n,t}) \le C_\gamma \gamma_{n,T}^{(0)} \le C_\gamma \gamma_{n,T},$$

where

$$C_\gamma = \frac{2}{\log\left(1+\sigma_{\mathrm{gp}}^{-2}\right)}.$$

Proof. Let

$$u_t = \sigma_{n,t-1}^2(\mathbf{x}_{n,t}).$$

Since conditioning cannot increase posterior variance and $k_n(\mathbf{x},\mathbf{x}) \le 1$ under Assumption B.2, we have $u_t \in [0,1]$. By concavity of $u \mapsto \log(1+\sigma_{\mathrm{gp}}^{-2}u)$ on $[0,1]$,

$$\log\left(1+\sigma_{\mathrm{gp}}^{-2}u_t\right) \ge u_t \log\left(1+\sigma_{\mathrm{gp}}^{-2}\right),$$

and hence

$$u_t \le \frac{\log\left(1+\sigma_{\mathrm{gp}}^{-2}u_t\right)}{\log\left(1+\sigma_{\mathrm{gp}}^{-2}\right)}.$$

For the realized BO query sequence, define the kernel matrix under the initial-data-conditioned covariance as

$$\mathbf{K}_{n,1:T}^{(0)} = \left[k_n^{(0)}(\mathbf{x}_{n,s},\mathbf{x}_{n,t})\right]_{s,t=1}^{T}.$$

The standard sequential determinant identity, now applied after conditioning on $\mathcal{D}_{n,0}$, gives

$$\frac{1}{2}\log\det\left(\mathbf{I}+\sigma_{\mathrm{gp}}^{-2}\mathbf{K}_{n,1:T}^{(0)}\right) = \frac{1}{2}\sum_{t=1}^{T}\log\left(1+\sigma_{\mathrm{gp}}^{-2}\sigma_{n,t-1}^2(\mathbf{x}_{n,t})\right).$$

Therefore,

$$\begin{aligned}\sum_{t=1}^{T}\sigma_{n,t-1}^2(\mathbf{x}_{n,t}) &\le \frac{2}{\log\left(1+\sigma_{\mathrm{gp}}^{-2}\right)}\frac{1}{2}\log\det\left(\mathbf{I}+\sigma_{\mathrm{gp}}^{-2}\mathbf{K}_{n,1:T}^{(0)}\right)\\ &\le C_\gamma\gamma_{n,T}^{(0)}\\ &\le C_\gamma\gamma_{n,T},\end{aligned}$$

where the final inequality follows from Eq. 4. □

B.3 Validity and Variance Scaling of FI-GP

Lemma B.3 (FI-GP covariance validity and variance scaling)**.** *Let the local GP posterior of agent n before round t be*

$$f_n(\mathbf{x}) \mid \mathcal{D}_{n,t-1} \sim \mathcal{GP}\left(\mu_{n,t-1}(\mathbf{x}), k_{n,t-1}(\mathbf{x},\mathbf{x}')\right).$$

For $\lambda_t \ge 0$, define

$$S_{n,t}(\mathbf{x}) = 1 + \lambda_t G_{n,t}(\mathbf{x})$$

and

$$\widetilde{k}_{n,t}(\mathbf{x},\mathbf{x}') = S_{n,t}(\mathbf{x})k_{n,t-1}(\mathbf{x},\mathbf{x}')S_{n,t}(\mathbf{x}').$$

Then $\widetilde{k}_{n,t}$ is a valid positive-semidefinite covariance kernel. Moreover, the FI-GP decision posterior preserves the local posterior mean and satisfies

$$\widetilde{\mu}_{n,t}(\mathbf{x}) = \mu_{n,t-1}(\mathbf{x}), \qquad \widetilde{\sigma}_{n,t}(\mathbf{x}) = S_{n,t}(\mathbf{x})\sigma_{n,t-1}(\mathbf{x}).$$

Proof. Consider an arbitrary finite set $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_q\}$. Let

$$\mathbf{K} = [k_{n,t-1}(\mathbf{x}_i, \mathbf{x}_j)]_{i,j=1}^{q}$$

and define

$$\mathbf{D}_S = \operatorname{diag}\left(S_{n,t}(\mathbf{x}_1), \dots, S_{n,t}(\mathbf{x}_q)\right).$$

The covariance matrix induced by $\widetilde{k}_{n,t}$ is

$$\widetilde{\mathbf{K}} = \mathbf{D}_S \mathbf{K} \mathbf{D}_S.$$

Since $\mathbf{K} \succeq 0$, for every vector $\mathbf{a} \in \mathbb{R}^q$,

$$\mathbf{a}^\top \widetilde{\mathbf{K}} \mathbf{a} = (\mathbf{D}_S \mathbf{a})^\top \mathbf{K} (\mathbf{D}_S \mathbf{a}) \geq 0.$$

Thus, $\widetilde{\mathbf{K}} \succeq 0$ for every finite set $\mathbf{X}$, and $\widetilde{k}_{n,t}$ is a valid covariance kernel.

The FI-GP mean is defined to equal the local posterior mean. For the marginal variance,

$$\begin{aligned}\widetilde{\sigma}^2_{n,t}(\mathbf{x}) &= \widetilde{k}_{n,t}(\mathbf{x}, \mathbf{x}) \\ &= S^2_{n,t}(\mathbf{x}) k_{n,t-1}(\mathbf{x}, \mathbf{x}) \\ &= S^2_{n,t}(\mathbf{x}) \sigma^2_{n,t-1}(\mathbf{x}).\end{aligned}$$

By Remark B.1 and $\lambda_t \geq 0$, we have $S_{n,t}(\mathbf{x}) \geq 1$. Taking square roots gives

$$\widetilde{\sigma}_{n,t}(\mathbf{x}) = S_{n,t}(\mathbf{x}) \sigma_{n,t-1}(\mathbf{x}).$$

□

B.4 Proof of Lemma B.1

Proof. Fix an arbitrary agent n and round t, and define

$$S_{n,t}(\mathbf{x}) = 1 + \lambda_t G_{n,t}(\mathbf{x}).$$

By the GUIDE-UCB selection rule in Eq. 1,

$$\mu_{n,t-1}(\mathbf{x}_{n,t}) + \sqrt{\beta_t} S_{n,t}(\mathbf{x}_{n,t}) \sigma_{n,t-1}(\mathbf{x}_{n,t}) \geq \mu_{n,t-1}(\mathbf{x}^*_n) + \sqrt{\beta_t} S_{n,t}(\mathbf{x}^*_n) \sigma_{n,t-1}(\mathbf{x}^*_n).$$

Rearranging gives

$$\begin{aligned}\mu_{n,t-1}(\mathbf{x}^*_n) - \mu_{n,t-1}(\mathbf{x}_{n,t}) \leq \sqrt{\beta_t} \Big[& S_{n,t}(\mathbf{x}_{n,t}) \sigma_{n,t-1}(\mathbf{x}_{n,t}) \\ & - S_{n,t}(\mathbf{x}^*_n) \sigma_{n,t-1}(\mathbf{x}^*_n) \Big].\end{aligned}$$

Under Assumption B.1,

$$f_n(\mathbf{x}^*_n) \leq \mu_{n,t-1}(\mathbf{x}^*_n) + \sqrt{\beta_t} \sigma_{n,t-1}(\mathbf{x}^*_n)$$

and

$$f_n(\mathbf{x}_{n,t}) \geq \mu_{n,t-1}(\mathbf{x}_{n,t}) - \sqrt{\beta_t} \sigma_{n,t-1}(\mathbf{x}_{n,t}).$$

Consequently,

$$\begin{aligned}r_{n,t} &= f_n(\mathbf{x}^*_n) - f_n(\mathbf{x}_{n,t}) \\ &\leq \mu_{n,t-1}(\mathbf{x}^*_n) - \mu_{n,t-1}(\mathbf{x}_{n,t}) + \sqrt{\beta_t} \sigma_{n,t-1}(\mathbf{x}^*_n) + \sqrt{\beta_t} \sigma_{n,t-1}(\mathbf{x}_{n,t}) \\ &\leq \sqrt{\beta_t} \Big[S_{n,t}(\mathbf{x}_{n,t}) \sigma_{n,t-1}(\mathbf{x}_{n,t}) - S_{n,t}(\mathbf{x}^*_n) \sigma_{n,t-1}(\mathbf{x}^*_n) \\ &\qquad + \sigma_{n,t-1}(\mathbf{x}^*_n) + \sigma_{n,t-1}(\mathbf{x}_{n,t}) \Big].\end{aligned}$$

Substituting $S_{n,t}(\mathbf{x}) = 1 + \lambda_t G_{n,t}(\mathbf{x})$ yields

$$\begin{aligned}r_{n,t} \leq \sqrt{\beta_t} \Big[& 2\sigma_{n,t-1}(\mathbf{x}_{n,t}) + \lambda_t G_{n,t}(\mathbf{x}_{n,t}) \sigma_{n,t-1}(\mathbf{x}_{n,t}) \\ & - \lambda_t G_{n,t}(\mathbf{x}^*_n) \sigma_{n,t-1}(\mathbf{x}^*_n) \Big],\end{aligned}$$

which proves Lemma B.1. □

B.5 Proof of Theorem 1

Proof. Fix an arbitrary agent n. By Lemma B.1, and because $\lambda_t \geq 0$ and $G_{n,t}(\mathbf{x}_n^*) \geq 0$, discarding the final nonpositive term gives

$$r_{n,t} \leq \sqrt{\beta_t}\,(2 + \lambda_t G_{n,t}(\mathbf{x}_{n,t}))\,\sigma_{n,t-1}(\mathbf{x}_{n,t}).$$

By Remark B.1, $G_{n,t}(\mathbf{x}_{n,t}) \leq G_{\max}$. Using $\lambda_t = \lambda_{\max}/\sqrt{t}$ with $\lambda_{\max} \geq 0$,

$$r_{n,t} \leq \sqrt{\beta_t}\left(2 + \frac{\lambda_{\max} G_{\max}}{\sqrt{t}}\right)\sigma_{n,t-1}(\mathbf{x}_{n,t}).$$

Since $\{\beta_t\}_{t\geq 1}$ is nondecreasing under Assumption B.1,

$$R_{n,T} \leq \sqrt{\beta_T}\sum_{t=1}^{T}\left(2 + \frac{\lambda_{\max} G_{\max}}{\sqrt{t}}\right)\sigma_{n,t-1}(\mathbf{x}_{n,t}).$$

Applying Cauchy–Schwarz to the complete product gives

$$R_{n,T} \leq \sqrt{\beta_T}\sqrt{\sum_{t=1}^{T}\left(2 + \frac{\lambda_{\max} G_{\max}}{\sqrt{t}}\right)^2} \times \sqrt{\sum_{t=1}^{T}\sigma_{n,t-1}^2(\mathbf{x}_{n,t})}.$$

Applying Lemma B.2 yields

$$R_{n,T} \leq \sqrt{\beta_T C_\gamma \gamma_{n,T}}\sqrt{\sum_{t=1}^{T}\left(2 + \frac{\lambda_{\max} G_{\max}}{\sqrt{t}}\right)^2}.$$

The remaining coefficient satisfies

$$\begin{aligned}\sum_{t=1}^{T}\left(2 + \frac{\lambda_{\max} G_{\max}}{\sqrt{t}}\right)^2 &= 4T + 4\lambda_{\max} G_{\max}\sum_{t=1}^{T} t^{-1/2} \\ &\quad + \lambda_{\max}^2 G_{\max}^2 \sum_{t=1}^{T} t^{-1} \\ &\leq 4T + 4\lambda_{\max} G_{\max}(2\sqrt{T} - 1) \\ &\quad + \lambda_{\max}^2 G_{\max}^2 (1 + \log T),\end{aligned}$$

where

$$\sum_{t=1}^{T} t^{-1/2} \leq 2\sqrt{T} - 1, \qquad \sum_{t=1}^{T} t^{-1} \leq 1 + \log T.$$

Combining these inequalities gives

$$R_{n,T} \leq \sqrt{\beta_T C_\gamma \gamma_{n,T}}\sqrt{4T + 4\lambda_{\max} G_{\max}(2\sqrt{T} - 1) + \lambda_{\max}^2 G_{\max}^2 (1 + \log T)}.$$

For fixed $\lambda_{\max}$ and $G_{\max}$, the guidance-dependent $\sqrt{T}$ and $\log T$ terms are lower order than the leading T term inside the second square root. Therefore,

$$R_{n,T} = \mathcal{O}\left(\sqrt{T\beta_T\gamma_{n,T}}\right).$$

□

B.6 Constant Intervention Strength

For a fixed intervention strength $\lambda_t \equiv \lambda$ with $\lambda \geq 0$, Lemma B.1 and Remark B.1 imply

$$r_{n,t} \leq \sqrt{\beta_t}(2 + \lambda G_{\max})\sigma_{n,t-1}(\mathbf{x}_{n,t}).$$

Using the monotonicity of β_t, Cauchy–Schwarz, and Lemma B.2,

$$\begin{aligned} R_{n,T} &\leq (2 + \lambda G_{\max})\sqrt{\beta_T}\sum_{t=1}^{T}\sigma_{n,t-1}(\mathbf{x}_{n,t}) \\ &\leq (2 + \lambda G_{\max})\sqrt{T\beta_T\sum_{t=1}^{T}\sigma^2_{n,t-1}(\mathbf{x}_{n,t})} \\ &\leq (2 + \lambda G_{\max})\sqrt{T\beta_T C_\gamma \gamma_{n,T}}. \end{aligned}$$

Hence, for any fixed $\lambda \geq 0$, the intervention changes only the constant factor in the regret bound and remains sublinear whenever $\beta_T\gamma_{n,T} = o(T)$.

B.7 Proof of Proposition 1

Proof. Fix an arbitrary agent n and round t. Let $\mathbf{x}_{n,t}$ be a selected suboptimal point, so that

$$\Delta_n(\mathbf{x}_{n,t}) = f_n(\mathbf{x}^*_n) - f_n(\mathbf{x}_{n,t}) = r_{n,t} > 0.$$

Starting from Lemma B.1,

$$\begin{aligned} \Delta_n(\mathbf{x}_{n,t}) \leq \sqrt{\beta_t}\Big[&2\sigma_{n,t-1}(\mathbf{x}_{n,t}) + \lambda_t G_{n,t}(\mathbf{x}_{n,t})\sigma_{n,t-1}(\mathbf{x}_{n,t}) \\ &- \lambda_t G_{n,t}(\mathbf{x}^*_n)\sigma_{n,t-1}(\mathbf{x}^*_n)\Big]. \end{aligned}$$

Rearranging gives

$$\sqrt{\beta_t}\,(2 + \lambda_t G_{n,t}(\mathbf{x}_{n,t}))\,\sigma_{n,t-1}(\mathbf{x}_{n,t}) \geq \Delta_n(\mathbf{x}_{n,t}) + \lambda_t\sqrt{\beta_t}G_{n,t}(\mathbf{x}^*_n)\sigma_{n,t-1}(\mathbf{x}^*_n).$$

Because $\beta_t > 0$ and $\lambda_t \geq 0$, the denominator below is strictly positive. Therefore,

$$\sigma_{n,t-1}(\mathbf{x}_{n,t}) \geq \frac{\Delta_n(\mathbf{x}_{n,t}) + \lambda_t\sqrt{\beta_t}G_{n,t}(\mathbf{x}^*_n)\sigma_{n,t-1}(\mathbf{x}^*_n)}{\sqrt{\beta_t}\,(2 + \lambda_t G_{n,t}(\mathbf{x}_{n,t}))}.$$

If $G_{n,t}(\mathbf{x}^*_n) \geq c$ and $G_{n,t}(\mathbf{x}_{n,t}) \leq \varepsilon_G$ for some $c > 0$ and $\varepsilon_G \geq 0$, then under $\lambda_t = \lambda_{\max}/\sqrt{t}$ with $\lambda_{\max} \geq 0$,

$$\sigma_{n,t-1}(\mathbf{x}_{n,t}) \geq \frac{\Delta_n(\mathbf{x}_{n,t}) + \frac{\lambda_{\max}c}{\sqrt{t}}\sqrt{\beta_t}\sigma_{n,t-1}(\mathbf{x}^*_n)}{\sqrt{\beta_t}\left(2 + \frac{\lambda_{\max}\varepsilon_G}{\sqrt{t}}\right)}.$$

This proves Proposition 1. □

Comparison with standard GP-UCB. Setting $\lambda_t = 0$ in Eq. 3 recovers the standard GP-UCB selection threshold $\Delta_n(\mathbf{x}_{n,t})/(2\sqrt{\beta_t})$. For $\lambda_t > 0$, direct comparison with the bound under $G_{n,t}(\mathbf{x}^*_n) \geq c$ and $G_{n,t}(\mathbf{x}_{n,t}) \leq \varepsilon_G$ shows that the GUIDE-UCB threshold is strictly larger whenever

$$2\sqrt{\beta_t}\,c\,\sigma_{n,t-1}(\mathbf{x}^*_n) > \varepsilon_G\Delta_n(\mathbf{x}_{n,t}).$$

Thus, the increase in the selection threshold occurs when guidance near the agent's optimum is sufficiently strong relative to the guidance assigned to the competing suboptimal point.

B.8 Geometric Interpretation of Informative Guidance

The informative-guidance condition admits a direct geometric interpretation. If a received component $p^\dagger$ lies within Mahalanobis radius ρ of the agent's optimum, then $G_{n,t}(\mathbf{x}^*_n) \geq \omega_{n,t,p^\dagger}\exp(-\rho^2/2)$. Conversely, if a suboptimal point $\mathbf{x}$ has Mahalanobis distance at least R_{tail} from every received component, then, because the retained weights sum to at most one, $G_{n,t}(\mathbf{x}) \leq \exp(-R^2_{\text{tail}}/2)$. Hence, the conditions $G_{n,t}(\mathbf{x}^*_n) \geq c$ and $G_{n,t}(\mathbf{x}) \leq \varepsilon_G$ arise naturally when the received guidance places substantial support near the agent's optimum while the suboptimal candidate lies in the tails of the received components.

C Benchmarks, Heterogeneity, Metrics, and Experimental Settings

C.1 Synthetic Benchmarks

Table 5 summarizes the 12 synthetic functions. We convert minimization functions to maximization by negating the function value where necessary. The implementation follows the corresponding BoTorch test-function implementations where available (Balandat et al., 2020), together with the Sphere, Weierstrass, Ellipsoid, and Zakharov implementations used in our code. The dimensions and search domains used in the experiments are reported explicitly in Table 5.

Table 5: Synthetic benchmark functions. All objectives are evaluated as maximization problems.

Function	**Dim.**	**Domain**	**Maximum**	**Landscape**
Ackley	10	$[-32.768, 32.768]^{10}$	0	Multimodal
Levy	10	$[-10, 10]^{10}$	0	Multimodal
Griewank	10	$[-600, 600]^{10}$	0	Multimodal
Rastrigin	10	$[-5.12, 5.12]^{10}$	0	Highly multimodal
Weierstrass	10	$[-0.5, 0.5]^{10}$	0	Irregular multimodal
Ellipsoid	10	$[-5.12, 5.12]^{10}$	0	Ill-conditioned
Sphere	10	$[-5, 5]^{10}$	0	Unimodal
Zakharov	10	$[-5, 5]^{10}$	0	Plate-shaped
Rosenbrock	10	$[-2.048, 2.048]^{10}$	0	Narrow curved valley
Michalewicz	10	$[0, \pi]^{10}$	9.66	Multimodal with steep valleys
Powell	10	$[-5, 5]^{10}$	0	Non-separable
Styblinski–Tang	10	$[-5, 5]^{10}$	391.66	Multimodal, separable

For Powell, we optimize a 10-dimensional search vector as implemented in our experiments; because the objective is evaluated in complete four-variable blocks, the first eight coordinates are active when $d = 10$, while the remaining two do not affect the function value.

All synthetic BO variables are represented in normalized search coordinates before local modeling and server-side component merging. The original domains in Table 5 define the underlying black-box functions, while the merging threshold δ_{merge} is applied in the normalized search space.

C.2 Real-world federated optimization tasks

Landmine Detection. The Landmine dataset (Xue et al., 2007) contains 29 binary classification tasks corresponding to 29 landmine fields, with each field treated as one FBO agent. For each field, the BO objective is evaluated using fixed three-fold stratified cross-validation on the local training data. Each queried configuration trains a support vector machine (SVM) with a radial basis function (RBF) kernel, with feature standardization fitted separately within each training fold. The optimized hyperparameters are

$$C_{\mathrm{SVM}} \in [10^{-4}, 10], \qquad \gamma_{\mathrm{RBF}} \in [10^{-2}, 10],$$

both decoded on logarithmic scales. The BO objective is the mean area under the receiver operating characteristic curve (AUC) across the cross-validation folds.

Activity Recognition Using Mobile Phone Sensors. Following FTS (Dai et al., 2020), we use the Human Activity Recognition Using Smartphones dataset from the University of California, Irvine (UCI) Machine Learning Repository. The dataset contains measurements from 30 subjects performing six activities, represented by 561 features. Each subject is treated as one FBO agent. The official partitions are merged, and the samples of each subject are divided into a stratified $50/50$ local training and validation split for each experimental replication. The features are further standardized using statistics from the local training split only.

Each BO query specifies three logistic regression hyperparameters:

$$B \in [20, 60], \qquad \lambda_{\mathrm{L2}} \in [10^{-6}, 1], \qquad \eta \in [10^{-2}, 10^{-1}],$$

where B is the mini-batch size, λ_{L2} is the L2 regularization parameter, and η is the learning rate. Each query trains a fresh linear classifier with 561 inputs and 6 outputs for 100 epochs using Stochastic Gradient Descent and cross-entropy loss. The BO objective is validation accuracy.

FedHPO. We use FedHPO-Bench (Wang et al., 2023b) to construct a federated hyperparameter optimization benchmark. Cora, CiteSeer, and PubMed are treated as three outer FBO agents, with each agent corresponding to a separate Graph Convolutional Network (GCN) hyperparameter optimization task under Federated Averaging (FedAvg). BO operates in a normalized search space, and each query is decoded according to the official FedHPO-Bench configuration space. Objective values are obtained from the official surrogate models using validation accuracy. We use full client participation and the highest available training-round fidelity.

C.3 Synthetic Task Heterogeneity and Empirical Similarity

Task construction. Generating related benchmark instances through transformations of a common base function is standard practice in black-box optimization. In particular, BBOB and COCO use search-space translations and rotations to generate different instances of benchmark functions (Hansen et al., 2009; 2021). Similar constructions have also been adopted in collaborative BO. For example, Chen et al. (2025) introduces agent heterogeneity by translating a common base function,

$$f_n(\mathbf{x}) = f_{\mathrm{base}}(\mathbf{x} + \boldsymbol{\psi}_n),$$

where $\boldsymbol{\psi}_n$ is sampled from a ball whose radius is proportional to the search-domain width.

Following this general principle, we combine agent-specific translations and rotations to control the degree of task heterogeneity. All variables are first represented in the normalized BO domain. For agent n, we define

$$f_n(\mathbf{x}) = f_{\mathrm{base}}\left(\mathbf{R}_n(\mathbf{x} - \mathbf{z}_n)\right),$$

where $\mathbf{z}_n$ controls the displacement of the landscape and $\mathbf{R}_n$ applies a rotation that, for non-rotationally-invariant functions, also changes the coordinate interactions.

Let $\mathbf{l}$ and $\mathbf{u}$ denote the lower and upper bounds of the normalized search domain. The shift vector is sampled as

$$\mathbf{z}_n \sim \mathcal{N}\left(\mathbf{0}, \left(\frac{\delta_{\mathrm{shift}}\|\mathbf{u} - \mathbf{l}\|_2}{\sqrt{d}}\right)^2 \mathbf{I}\right).$$

Hence,

$$\sqrt{\mathbb{E}\|\mathbf{z}_n\|_2^2} = \delta_{\mathrm{shift}}\|\mathbf{u} - \mathbf{l}\|_2,$$

so δ_{shift} directly controls the root-mean-square shift magnitude relative to the domain scale.

For the rotation, we draw $\mathbf{A}_n \in \mathbb{R}^{d\times d}$ with independent standard Gaussian entries and construct the skew-symmetric matrix

$$\mathbf{H}_n = \mathbf{A}_n - \mathbf{A}_n^\top.$$

We then define

$$\mathbf{R}_n = \exp\left(\delta_{\mathrm{rot}}\mathbf{H}_n\right).$$

Since the matrix exponential of a skew-symmetric matrix is orthogonal, $\mathbf{R}_n$ is a valid rotation matrix. The parameter δ_{rot} controls the strength of the rotation.

We consider three regimes:

$$\begin{aligned}
&\text{Level 1:} \quad (\delta_{\mathrm{shift}}, \delta_{\mathrm{rot}}) = (0, 0),\\
&\text{Level 2:} \quad (\delta_{\mathrm{shift}}, \delta_{\mathrm{rot}}) = (0.05, 0.1),\\
&\text{Level 3:} \quad (\delta_{\mathrm{shift}}, \delta_{\mathrm{rot}}) = (0.3, 1.0).
\end{aligned}$$

Level 1 contains identical local objectives, Level 2 introduces mild heterogeneity, and Level 3 applies substantially stronger transformations. For every experimental replication, the same agent-specific transformations are used by all compared methods.

Empirical task similarity. The transformation parameters specify how tasks are generated but do not directly quantify the similarity of the resulting objective landscapes. We therefore complement them with a rank-based empirical similarity measure.

For each benchmark b, heterogeneity level ℓ, and experimental replication r, we generate a common scrambled Sobol probe set

$$\mathcal{X}_{\text{probe}} = \{\mathbf{x}_q\}_{q=1}^{Q}, \qquad Q = 4096,$$

over the normalized search domain using a fixed Sobol seed of zero. For agent n, we evaluate its noise-free latent objective on all probe points and form

$$\mathbf{v}_n^{(b,\ell,r)} = [f_n(\mathbf{x}_1), \ldots, f_n(\mathbf{x}_Q)]\,.$$

For every pair of agents $n < m$, their similarity is measured by the Spearman rank correlation

$$\rho_{n,m}^{(b,\ell,r)} = \rho_{\text{S}}\left(\mathbf{v}_n^{(b,\ell,r)}, \mathbf{v}_m^{(b,\ell,r)}\right).$$

The mean similarity of one benchmark realization is

$$\text{Sim}_{b,\ell,r} = \frac{2}{N(N-1)} \sum_{n<m} \rho_{n,m}^{(b,\ell,r)}.$$

Finally, we average over the $B = 12$ synthetic benchmarks and $R = 10$ experimental realizations:

$$\overline{\text{Sim}}_{\ell} = \frac{1}{BR} \sum_{b=1}^{B} \sum_{r=1}^{R} \text{Sim}_{b,\ell,r}.$$

Spearman correlation measures agreement in the relative ranking of common candidate locations and is insensitive to differences in objective scale. We use it only to characterize the generated task heterogeneity.

Table 6: Mean pairwise Spearman correlation between the 16 agent objectives. Each benchmark entry is averaged over 10 task realizations, and Overall additionally averages across the 12 benchmarks. Higher values indicate greater agreement in candidate rankings.

Benchmark	**Level 1**	**Level 2**	**Level 3**
Ackley	1.000	0.485	0.027
Levy	1.000	0.450	0.076
Griewank	1.000	0.861	0.155
Rastrigin	1.000	0.489	0.136
Weierstrass	1.000	0.464	0.005
Ellipsoid	1.000	0.740	0.120
Sphere	1.000	0.861	0.155
Zakharov	1.000	0.527	0.013
Rosenbrock	1.000	0.676	0.115
Michalewicz	1.000	0.034	-0.001
Powell	1.000	0.635	0.046
Styblinski–Tang	1.000	0.466	0.124
Overall	**1.000**	**0.557**	**0.081**

The empirical similarities clearly separate the three heterogeneity regimes. Level 1 is exactly homogeneous, Level 2 retains moderate agreement in candidate rankings, and Level 3 exhibits only weak average agreement across agents.

C.4 Evaluation Metrics

For a synthetic objective, let $f_n^* = f_n(\mathbf{x}_n^*)$ denote the optimal value of agent n, and let $f_{n,t}^{\text{best}}$ denote its best noise-free value observed through round t. The simple regret is

$$s_{n,t} = f_n^* - f_{n,t}^{\text{best}},$$

and the reported average simple regret is

$$\overline{s}_t = \frac{1}{N}\sum_{n=1}^{N} s_{n,t}.$$

Lower values are better.

For Landmine Detection, the reported performance is AUC averaged across the 29 agents. For Activity Recognition and FedHPO, the reported performance is validation accuracy averaged across their agents. Higher values are better for these real-world metrics. For convergence plots, we use the lower-is-better metrics $1-\mathrm{AUC}$ and $1-\mathrm{Accuracy}$. The main text reports the original AUC and accuracy values.

Average ranks are computed independently within each benchmark and then averaged over the corresponding benchmark set.

C.5 Shared Experimental Settings

All compared methods use identical search domains, per-agent evaluation budgets, and repetition seeds within each benchmark. For a fair comparison, all baseline methods use the hyperparameter settings recommended in their original papers or official implementations. For synthetic experiments, all methods except FTS-DE use the same Collaborative Latin Hypercube Sampling (CLHS) initialization following Chen et al. (2025). FTS-DE retains the initialization prescribed by its Distributed Exploration mechanism (Dai et al., 2021), which assigns agents to different search subregions during initialization. Synthetic experiments use $N = 16$ agents, 30 initial observations per agent, 50 BO rounds, and Gaussian observation noise with standard deviation 0.1. The real-world experiments use their natural task counts: 29 agents for Landmine Detection, 30 agents for Activity Recognition, and 3 agents for FedHPO. Each experiment uses 30 initial evaluations per agent and 50 BO rounds, and the complete federated experiment is repeated independently 10 times.

GUIDE-FBO configuration. GUIDE-FBO introduces four principal tunable hyperparameters: the number of posterior-optimum samples M, the downlink packet size P, the merging threshold δ_{merge}, and the maximum guidance strength $\lambda_{\max}$. Unless otherwise stated, the same configuration is used across all benchmarks without benchmark-specific tuning. The remaining quantities in Table 7 are fixed modeling or implementation choices.

Table 7: Default configuration of GUIDE-FBO. The same settings are used across benchmarks unless otherwise stated.

Parameter	**Setting**
Posterior-optimum samples M	500
ORF features D_{ORF}	500
DPGMM covariance	Diagonal
LCB coefficient κ	1.0
Merge threshold δ_{merge}	0.05
Downlink packet size P	5
Maximum guidance strength $\lambda_{\max}$	1.0
Guidance schedule λ_t	$\lambda_{\max}/\sqrt{t}$

The server-side softmax temperature is determined automatically from the spread of the standardized component scores rather than treated as an independently tuned hyperparameter. Specifically, letting

$$\Delta v_t = \max_{\ell} v'_\ell - \min_{\ell} v'_\ell,$$

we use

$$\tau_t = \begin{cases} \max\{0.1\Delta v_t,\, 10^{-2}\}, & \Delta v_t > 10^{-6}, \\ 10^{-2}, & \text{otherwise.} \end{cases}$$

The local surrogate is implemented as a fixed-noise GP with an ARD Matérn-$5/2$ kernel, whose hyperparameters are fitted by maximizing the exact marginal likelihood. The ORF posterior-path approximation follows the corresponding Matérn-$5/2$ spectral construction described in Appendix I.

D Complete Experimental Results and Convergence Trajectories

This section reports the complete numerical results and convergence trajectories underlying Section 5.2. For each synthetic heterogeneity level, we report the final average simple regret on all 12 benchmark functions and the corresponding trajectories over the 50 BO rounds. We then report the trajectories for the three real-world tasks. All curves show the mean over 10 independent runs, with shaded regions denoting ± 1 standard error of the mean.

D.1 Level 1: $\delta_{\text{shift}} = 0$ and $\delta_{\text{rot}} = 0$

Table 8: Complete synthetic results under Level 1 (Homogeneous). Entries report the final mean average simple regret over 10 independent runs. Lower is better. The best and second-best results in each row are shown in bold and underlined, respectively.

Benchmark	TS	UCB	NEI	FTS	FTS-DE	FMTBO	CGP-TS	CGP-UCB	CGP-NEI	GUIDE-TS	GUIDE-UCB	GUIDE-NEI
Ackley	19.4789	19.3082	18.4994	19.4870	19.6941	19.0814	19.5242	8.7013	6.5675	18.6182	**4.4896**	5.2015
Levy	24.9042	11.9439	10.5159	23.0893	23.5361	9.3873	22.9011	3.7128	2.6747	11.6750	**1.5865**	1.9122
Griewank	47.1321	14.1117	13.3126	43.9996	15.4063	11.3637	43.2063	6.3414	4.9622	23.2003	**4.9517**	6.3052
Rastrigin	105.6477	80.4791	79.9759	102.8319	103.4912	81.3840	102.9572	62.0303	59.7812	90.3331	**53.4890**	54.6122
Weierstrass	13.5724	12.8962	11.8821	13.5131	13.5052	12.2728	13.5183	8.0722	8.0706	12.4086	**3.2670**	3.7779
Ellipsoid	146.4074	117.5092	87.9959	138.8153	137.5396	81.6766	140.1417	26.8736	24.0600	98.2097	**1.7311**	3.0529
Sphere	23.9937	14.8929	11.2488	23.4224	21.1340	9.8874	22.6427	3.9151	2.9221	11.6519	**0.6487**	1.0040
Zakharov	60.3082	47.1536	46.0734	58.7893	60.6199	46.2480	59.6063	48.3546	45.0008	44.7664	**15.6653**	22.4356
Rosenbrock	376.8339	110.0299	105.3660	363.1293	270.5447	111.8808	340.1700	90.6186	73.0144	162.5873	**45.4639**	51.9817
Michalewicz	6.3605	5.8237	5.8773	6.4475	6.2028	6.2124	6.4715	5.2263	**5.0059**	6.3292	5.2700	5.2798
Powell	677.0909	139.5000	123.8940	716.2145	707.2955	116.0112	651.4241	195.2524	116.9121	329.9346	**38.6966**	38.7592
Styblinski–Tang	156.5159	107.2154	103.6553	156.2315	155.8989	103.3376	153.1554	62.7248	**56.2330**	150.9906	99.6545	102.8457

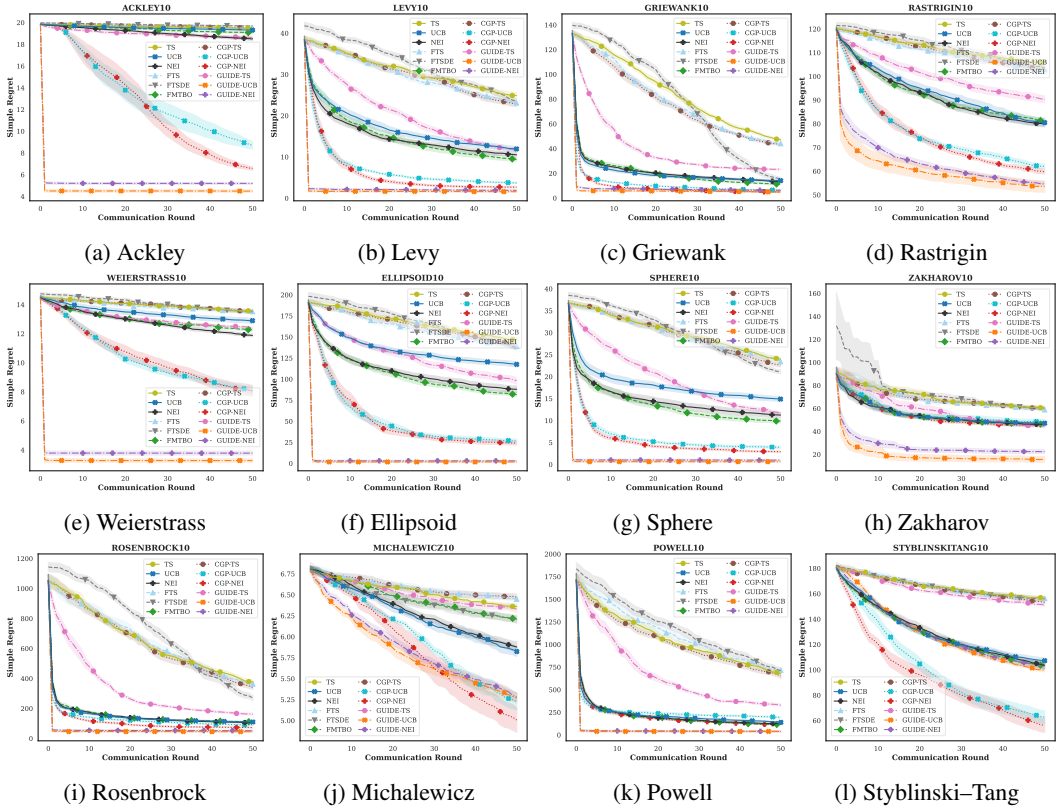

(a) Ackley (b) Levy (c) Griewank (d) Rastrigin

(e) Weierstrass (f) Ellipsoid (g) Sphere (h) Zakharov

(i) Rosenbrock (j) Michalewicz (k) Powell (l) Styblinski–Tang

Figure 4: Convergence trajectories for all 12 synthetic functions under Level 1 (Homogeneous). Curves show the mean average simple regret over 10 independent runs, and shaded regions denote ± 1 standard error of the mean across runs. Lower is better.

Level 1 represents the fully aligned setting. All agents optimize the same latent objective, and the empirical mean pairwise Spearman correlation is 1.000. GUIDE-UCB achieves the best average rank of 1.33 and obtains the lowest final average simple regret on 10 of the 12 benchmarks. GUIDE-NEI ranks second overall with an average rank of 2.42 and is among the two best methods on nine benchmarks. Moreover, GUIDE-UCB, GUIDE-NEI, and GUIDE-TS each outperform their matched independent acquisition rule on all 12 functions.

The numerical gaps are also large on several benchmarks. UCB and GUIDE-UCB obtain final regrets of 117.5092 and 1.7311 on Ellipsoid, 14.8929 and 0.6487 on Sphere, 47.1536 and 15.6653 on Zakharov, 110.0299 and 45.4639 on Rosenbrock, and 139.5000 and 38.6966 on Powell. GUIDE-UCB also obtains 4.4896 on Ackley and 1.5865 on Levy, compared with 8.7013 and 3.7128 for CGP-UCB. These results show that, when the local objectives are fully aligned, exchanging distributions over the locations of the agents' optima can substantially reduce simple regret within a fixed evaluation budget.

The convergence trajectories in Figure 4 make this effect more apparent. Except on Michalewicz and Styblinski–Tang, GUIDE-UCB and GUIDE-NEI separate from most competing methods relatively early and remain in the lowest-regret group over much of the subsequent optimization horizon. This behavior is consistent with the mechanism of GUIDE. Because the agents optimize the same objective, their uploaded components are more likely to concentrate around the same promising regions. Server-side merging can therefore combine consistent information from different agents. FI-GP then increases local posterior uncertainty selectively around these regions without changing the local posterior mean. The resulting FI-GP makes these globally supported regions more attractive to UCB and NEI when they remain plausible under the local posterior, which can lead them to be evaluated earlier.

GUIDE-TS shows a smaller visual advantage. Unlike UCB and NEI, TS selects each query by maximizing a random sample path drawn from FI-GP. Federated guidance changes the distribution from which this path is sampled, but the realized path can still attain its maximum elsewhere. This additional sampling variability makes the effect of spatial guidance less consistent across rounds. GUIDE-TS still improves upon independent TS on all 12 functions, but does not show the same clear early separation as GUIDE-UCB and GUIDE-NEI.

Michalewicz and Styblinski–Tang are the two main exceptions. Michalewicz contains multiple steep valleys, so its high-quality regions are spatially narrow. Region-level guidance is less useful when its spatial support does not closely match the relevant valley. Since the tasks are identical at Level 1, the high-potential points transferred by CGP are directly relevant to every agent, which provides a favorable setting for point-level transfer. CGP-NEI and CGP-UCB accordingly obtain 5.0059 and 5.2263, compared with 5.2700 and 5.2798 for GUIDE-UCB and GUIDE-NEI.

Styblinski–Tang has many competing local basins with similar structure. Because several local basins can remain plausible, a distribution over optimum locations may provide a less precise signal than a directly transferred high-quality point. Under Level 1, where all agents optimize exactly the same objective, directly transferring a good point can be more efficient because the same point is useful to every agent. This favors the point-level transfer used by CGP. CGP-NEI and CGP-UCB obtain final regrets of 56.2330 and 62.7248, while GUIDE-UCB and GUIDE-NEI obtain 99.6545 and 102.8457. These two exceptions show that, even when the tasks are identical, the most effective form of transferred information still depends on the geometry of the objective.

D.2 Level 2: $\delta_{\text{shift}} = 0.05$ and $\delta_{\text{rot}} = 0.1$

Level 2 introduces mild task heterogeneity, with the empirical mean pairwise Spearman correlation decreasing from 1.000 to 0.557. GUIDE-UCB remains the strongest aggregate method with an average rank of 1.58. It achieves the lowest final regret on 9 benchmarks and ranks among the top two methods on 11 of the 12 functions. GUIDE-NEI ranks second overall at 2.17, is the best method on Zakharov and Powell, and appears among the top two methods on 10 functions. Relative to their matched independent baselines, GUIDE-UCB and GUIDE-NEI improve on 11 of the 12 functions, while GUIDE-TS improves on all 12.

Large absolute differences remain on several benchmarks. UCB and GUIDE-UCB obtain 19.4350 and 8.4785 on Ackley, 113.6023 and 13.1331 on Ellipsoid, 14.2443 and 2.2689 on Sphere, and

Table 9: Complete synthetic results under Level 2 (Mild heterogeneity). Entries report the final mean average simple regret over 10 independent runs. Lower is better. The best and second-best results in each row are shown in bold and underlined, respectively.

Benchmark	TS	UCB	NEI	FTS	FTS-DE	FMTBO	CGP-TS	CGP-UCB	CGP-NEI	GUIDE-TS	GUIDE-UCB	GUIDE-NEI
Ackley	19.5136	19.4350	18.7628	19.5511	19.7171	18.9928	19.5347	10.1460	10.9368	18.8906	**8.4785**	9.5379
Levy	24.3585	11.0529	10.1363	23.4060	25.8641	9.7791	24.2127	4.3481	3.5741	12.7145	**3.2816**	3.3482
Griewank	44.4283	13.5953	11.9247	44.1421	26.2381	11.3470	43.6216	7.1888	7.1002	24.8356	**6.3000**	7.8444
Rastrigin	101.5417	80.4771	78.9030	101.4711	101.7745	80.9836	103.0048	64.6460	62.5459	88.0868	**58.3082**	60.1472
Weierstrass	13.4680	12.9344	12.1369	13.4517	13.6576	12.5151	13.5531	8.9999	9.3717	12.7340	**6.6966**	6.8316
Ellipsoid	139.8685	113.6023	86.7074	139.5091	130.8015	78.4957	136.6720	37.2115	29.1808	91.6564	**13.1331**	13.3723
Sphere	21.6883	14.2443	10.2896	23.9013	20.9801	9.0656	21.6460	4.5267	3.3859	11.3485	**2.2689**	2.5063
Zakharov	62.9710	49.8750	48.2225	62.7988	62.9958	47.6923	62.7797	50.1082	47.1746	48.4994	22.1717	**21.8975**
Rosenbrock	385.1965	97.8257	104.8257	365.1723	356.7614	107.8170	392.6816	98.7769	83.6536	183.4833	**55.2114**	59.3597
Michalewicz	6.6497	6.1248	6.1995	6.5934	6.5978	6.4469	6.6054	6.1800	6.1802	6.5636	**5.8629**	5.8783
Powell	701.0029	157.3605	126.0791	746.8994	853.3029	117.3149	710.5802	205.5435	116.5086	348.1294	94.6703	**74.3549**
Styblinski–Tang	159.3047	114.4204	114.9073	160.0953	162.8046	119.9616	157.4158	127.6079	**109.7452**	158.4985	123.2241	116.9887

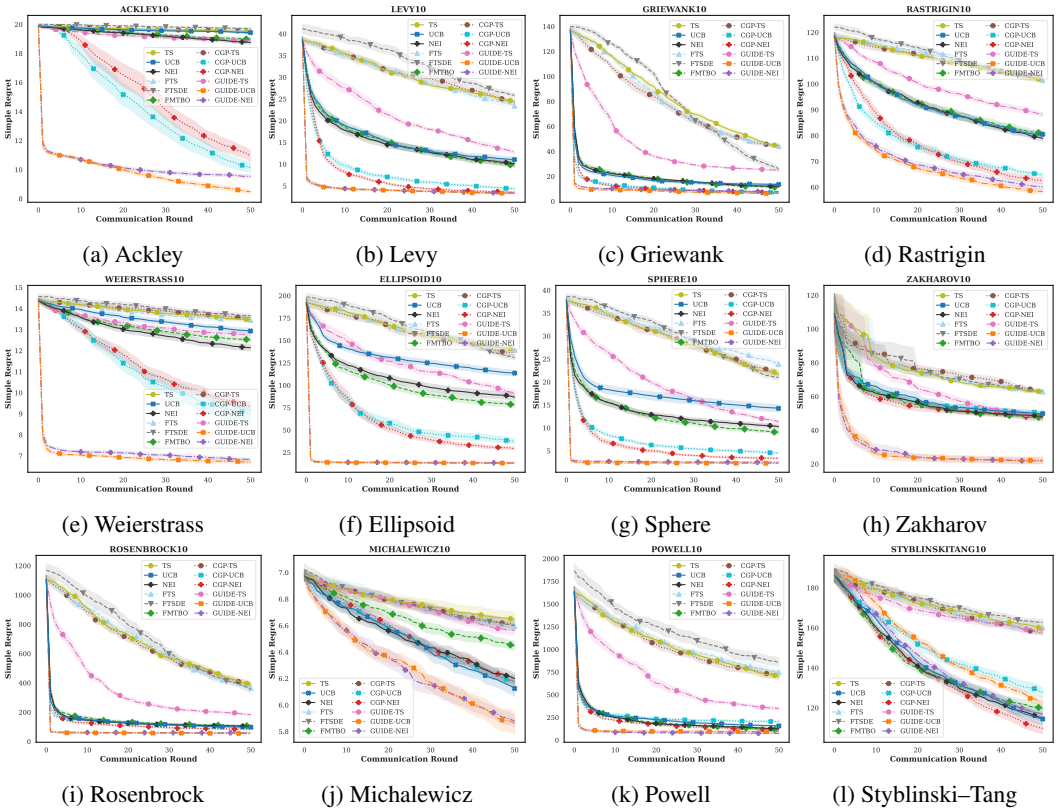

(a) Ackley (b) Levy (c) Griewank (d) Rastrigin

(e) Weierstrass (f) Ellipsoid (g) Sphere (h) Zakharov

(i) Rosenbrock (j) Michalewicz (k) Powell (l) Styblinski–Tang

Figure 5: Convergence trajectories for all 12 synthetic functions under Level 2 (Mild heterogeneity). Curves show the mean average simple regret over 10 independent runs, and shaded regions denote ± 1 standard error of the mean across runs. Lower is better.

97.8257 and 55.2114 on Rosenbrock. On Zakharov, GUIDE-NEI obtains 21.8975, compared with 47.1746 for CGP-NEI and 47.6923 for FMTBO. On Powell, GUIDE-NEI obtains 74.3549, compared with 116.5086 for CGP-NEI and 117.3149 for FMTBO. The benefit of exchanging distributions over the locations of the agents' optima therefore remains substantial after moderate shifts and rotations are introduced.

The trajectories in Figure 5 retain much of the qualitative pattern observed at Level 1. GUIDE-UCB and GUIDE-NEI show a particularly clear lead on Ackley, Levy, Weierstrass, Ellipsoid, Sphere, and Zakharov, where their curves separate visibly from the main competing methods during the optimization process. On Griewank, Rastrigin, Rosenbrock, Michalewicz, and Powell, the two GUIDE variants also reach competitive or leading low-regret regions, but the gap from the strongest baselines is smaller. Styblinski–Tang remains the main exception. Compared with Level 1, GUIDE

still reaches low-regret regions earlier on most benchmarks, but its lead over the strongest baselines becomes smaller on several functions.

This difference follows directly from the task construction. At Level 1, the distributions transferred by different agents refer to the same underlying objective and therefore tend to support the same promising regions. At Level 2, the shifts move these regions and the rotations change the local geometry of non-rotationally-invariant objectives. A promising region for one agent may therefore no longer align exactly with the corresponding region of another agent. Since these transformations are still moderate, useful spatial structure can remain shared across agents. GUIDE transfers distributions over promising regions, so a small spatial mismatch does not invalidate the shared information as easily as it can invalidate a single transferred design.

The change on Michalewicz is especially informative. Under Level 1, CGP-NEI and CGP-UCB outperform the GUIDE variants, whereas under Level 2 GUIDE-UCB and GUIDE-NEI obtain the two lowest final regrets of 5.8629 and 5.8783. A likely explanation is that point-level transfer is most effective when the narrow high-quality valleys of different agents are closely aligned. Once mild shifts and rotations are introduced, a high-quality design from one agent can fall outside the corresponding valley of another agent. GUIDE instead transfers a distribution over a promising region, making the transferred information less dependent on exact point alignment. This broader spatial representation is therefore better suited to the mild mismatch introduced at Level 2.

Michalewicz also illustrates the distinction between global task similarity and local transferability. Its Level 2 mean pairwise Spearman correlation is only 0.034, yet GUIDE-UCB and GUIDE-NEI obtain the two best final results. The Spearman correlation measures ranking agreement over the full search space, whereas GUIDE acts on the spatial support of promising regions. A low global correlation can therefore coexist with useful shared structure near high-value regions.

Styblinski–Tang remains the main exception. Its many competing local basins make it difficult for a regional distribution to indicate one clearly preferred search direction. Under Level 1, point-level transfer is especially effective because a good point identified by one agent is directly useful to all others. After the Level 2 transformations are introduced, this exact point correspondence becomes weaker. This is visible in the final results. CGP-UCB changes from 62.7248 at Level 1 to 127.6079 at Level 2, while CGP-NEI changes from 56.2330 to 109.7452. The corresponding GUIDE-UCB results are 99.6545 and 123.2241, and the GUIDE-NEI results are 102.8457 and 116.9887. GUIDE is therefore still not the best method on Styblinski–Tang at Level 2, but its performance deteriorates much less than the two CGP variants as mild heterogeneity is introduced. On this benchmark, the smaller degradation of GUIDE suggests that distributional guidance is less sensitive to mild spatial misalignment than point-level transfer.

D.3 Level 3: $\delta_{\text{shift}} = 0.3$ and $\delta_{\text{rot}} = 1.0$

Table 10: Complete synthetic results under Level 3 (Severe heterogeneity). Entries report the final mean average simple regret over 10 independent runs. Lower is better. The best and second-best results in each row are shown in bold and underlined, respectively.

Benchmark	TS	UCB	NEI	FTS	FTS-DE	FMTBO	CGP-TS	CGP-UCB	CGP-NEI	GUIDE-TS	GUIDE-UCB	GUIDE-NEI
Ackley	20.4047	20.5073	20.4203	20.4080	20.3972	20.2383	20.3699	20.3380	20.4331	20.3552	20.1739	**20.1161**
Levy	37.5643	24.3700	25.5279	41.1159	40.0354	25.0598	34.5967	24.6258	30.4253	34.2904	**20.8306**	23.6984
Griewank	67.8144	28.8836	26.9351	72.7054	74.5664	27.0838	65.7857	21.7611	30.0882	68.9316	**21.2957**	29.7074
Rastrigin	106.3844	95.6377	96.0487	110.2650	111.8623	95.4028	108.3831	96.2275	100.1082	104.5270	**90.8849**	94.0912
Weierstrass	13.6349	13.1439	12.7244	13.6075	13.6386	12.8031	13.4763	12.9093	12.8169	13.4107	12.8443	**12.4937**
Ellipsoid	157.2353	162.9886	138.9518	172.5919	167.3659	**120.5134**	156.5689	134.2802	151.8904	149.4864	127.6623	127.9452
Sphere	24.1755	22.4991	19.2331	26.1404	26.1197	**15.6632**	22.9534	18.3412	21.1273	24.0913	17.6787	17.9497
Zakharov	196.5950	101.9986	1061.0835	1095.3250	139.9111	542.0744	167.0674	103.6758	601.0594	128.2162	**92.9444**	307.8536
Rosenbrock	911.4690	**257.8472**	316.0209	1036.4767	1056.0210	294.0791	873.7124	299.9545	457.7168	919.0478	321.9631	331.0628
Michalewicz	7.6120	**7.3490**	7.4127	7.6202	7.6026	7.4742	7.6412	7.3963	7.3846	7.6373	7.3793	7.4323
Powell	1397.4194	681.4695	793.8595	1683.7471	1659.8198	751.6538	1267.8210	718.8015	1046.9279	1419.2525	685.3075	**656.6218**
Styblinski–Tang	199.0988	157.1972	149.6469	206.1874	207.4384	**148.3097**	193.8517	169.4907	157.8670	196.6009	165.9432	149.2083

Level 3 introduces much stronger task heterogeneity. The larger translations move the promising regions of different agents farther apart, while the stronger rotations can substantially change the geometry of non-rotationally-invariant objectives. The empirical mean pairwise Spearman correlation decreases to 0.081. The agents therefore share much less directly transferable spatial information than under Levels 1 and 2.

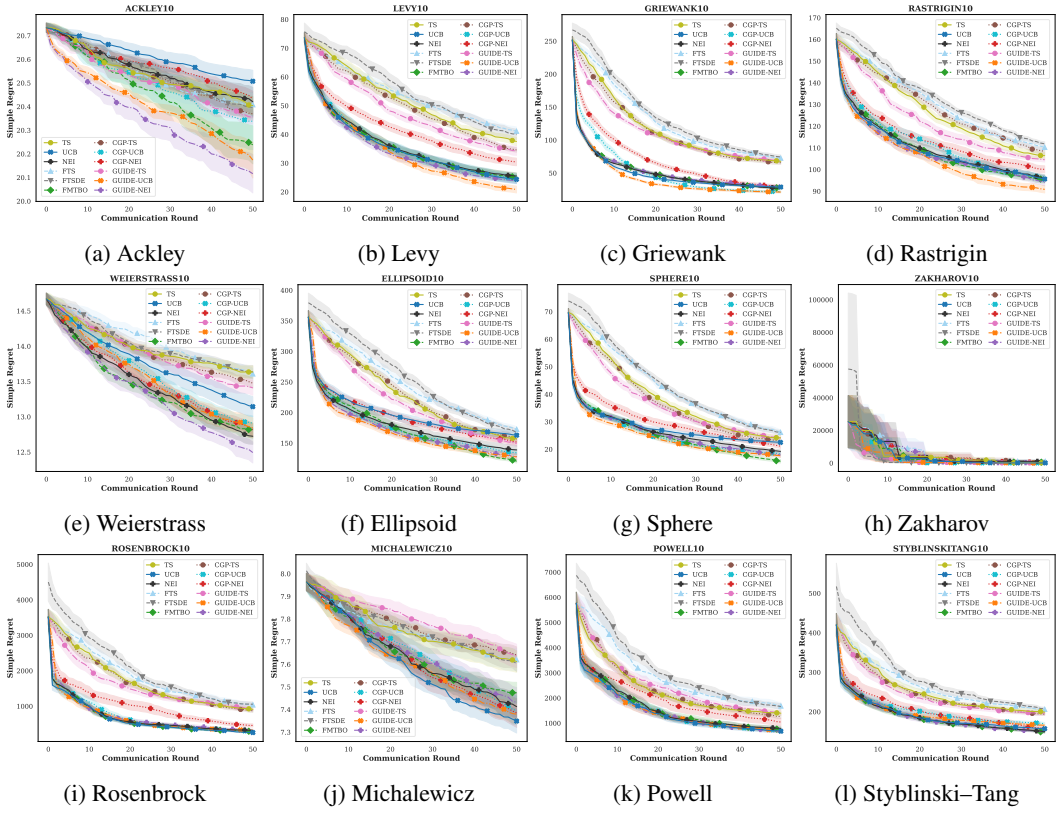

(a) Ackley (b) Levy (c) Griewank (d) Rastrigin

(e) Weierstrass (f) Ellipsoid (g) Sphere (h) Zakharov

(i) Rosenbrock (j) Michalewicz (k) Powell (l) Styblinski–Tang

Figure 6: Convergence trajectories for all 12 synthetic functions under Level 3 (Severe heterogeneity). Curves show the mean average simple regret over 10 independent runs, and shaded regions denote ± 1 standard error of the mean across runs. Lower is better.

This setting is particularly difficult for point-level knowledge transfer. A promising design found by one agent may no longer lie in a promising region for another agent after the stronger shifts and rotations. The CGP results illustrate this limitation. CGP-TS and CGP-UCB only slightly improve their average ranks over independent TS and UCB, from 9.17 to 8.33 and from 4.83 to 4.25, respectively. In contrast, CGP-NEI has an average rank of 6.75, worse than 5.42 for independent NEI. Point-level transfer therefore no longer provides a consistent advantage when the agents' optima are poorly aligned.

FMTBO becomes more competitive in this regime and achieves an average rank of 3.67, the best result among the non-GUIDE methods. It is also the best method on Ellipsoid, Sphere, and Styblinski–Tang. Unlike methods that transfer particular promising locations, FMTBO communicates model-level information through GP hyperparameters and estimates task relatedness from predictive rankings. Such information does not require the optimum of one agent to occur at nearly the same location as the optimum of another agent. This weaker dependence on exact spatial alignment helps explain why FMTBO becomes relatively more competitive at Level 3.

Despite the severe heterogeneity, GUIDE-UCB retains the best aggregate rank of 2.58, while GUIDE-NEI ranks second at 3.42. GUIDE-UCB is among the top two methods on 8 of the 12 benchmarks and obtains the lowest final regret on Levy, Griewank, Rastrigin, and Zakharov. GUIDE-NEI is best on Ackley, Weierstrass, and Powell. The complete trajectories in Figure 6 show a much less uniform advantage than at Levels 1 and 2. The leading curves overlap and cross more frequently, and GUIDE no longer separates early from the other methods on most benchmarks. This is expected because the received distributions now contain more information that is useful to some agents but less relevant to others.

Two mechanisms help GUIDE remain robust in this regime. First, the strength of federated guidance decays according to $\lambda_t = \lambda_{\max}/\sqrt{t}$. Global information therefore has its largest influence when

local observations are scarce, while its influence gradually decreases as each agent collects more evidence about its own objective. The direct influence of global guidance therefore decreases over rounds, allowing local evidence to play a larger role later in the optimization.

Second, FI-GP leaves the local posterior mean unchanged and rescales only the posterior uncertainty. The absolute uncertainty increment at $\mathbf{x}$ is $\lambda_t G_{n,t}(\mathbf{x})\sigma_{n,t-1}(\mathbf{x})$. Global guidance therefore has a large effect mainly where the federation provides strong support and the local GP remains uncertain. If a region has already been sufficiently explored locally, its posterior uncertainty is typically smaller, so even strong global guidance produces only a limited intervention. If local evaluations also indicate poor performance, the low posterior mean is preserved. Such a region is therefore unlikely to be selected solely because it is globally recommended.

The remaining failures also follow this interpretation. On Rosenbrock, independent UCB obtains 257.8472, compared with 321.9631 for GUIDE-UCB. Its narrow curved valley makes transferred spatial guidance particularly sensitive to strong translations and rotations. On Michalewicz, UCB obtains 7.3490, while GUIDE-UCB obtains 7.3793. The difference is much smaller, but the steep and narrow high-quality regions again leave little tolerance for spatial mismatch. These cases show that negative transfer can still occur on individual tasks. The main Level 3 result is that GUIDE retains the best aggregate rank despite these task-specific failures.

D.4 Real-World Convergence Trajectories

Figure 7 reports the complete convergence trajectories corresponding to Table 2. For visualization, we use $1 - \mathrm{AUC}$ for Landmine Detection and $1 - \mathrm{Accuracy}$ for Activity Recognition and FedHPO, so lower values are better.

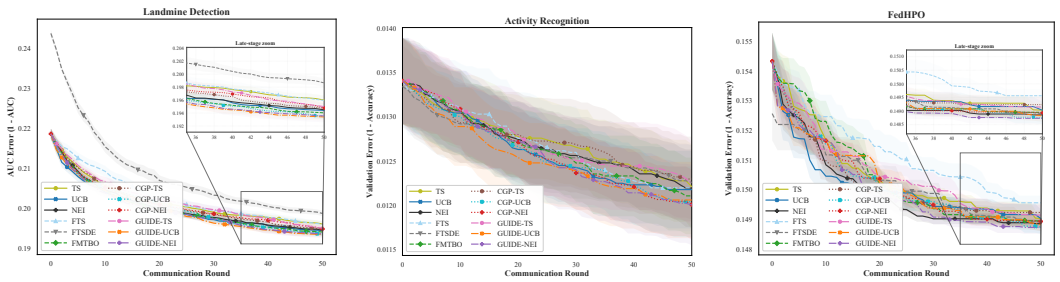

(a) Landmine Detection (b) Activity Recognition (c) FedHPO

Figure 7: Convergence trajectories on the three real-world federated optimization tasks. Curves show the mean transformed error metric over 10 independent runs, and shaded regions denote ± 1 standard error of the mean. Lower is better.

On the three real-world tasks, the strongest methods often finish within a narrow performance range, while the matched comparisons remain consistently favorable for GUIDE-UCB and GUIDE-NEI.

On Landmine Detection, GUIDE-UCB achieves the highest final mean AUC of 0.806604, followed by CGP-UCB at 0.806501 and GUIDE-NEI at 0.806452. Independent UCB obtains 0.805593. The convergence trajectories remain close during the earlier rounds, while GUIDE-UCB gradually moves to the best mean performance later in the optimization horizon. Unlike many Level 1 synthetic benchmarks, there is no clear early separation from the strongest baselines.

Activity Recognition is much more tightly clustered. GUIDE-NEI obtains the highest final mean validation accuracy of 0.987994, followed by CGP-NEI at 0.987990 and CGP-UCB at 0.987959. GUIDE-UCB obtains 0.987938, compared with 0.987821 for independent UCB. All leading methods remain within a narrow accuracy range over most of the BO horizon. This task therefore shows that GUIDE remains competitive, but does not exhibit a large practical separation among the strongest methods under the tested budget.

FedHPO shows a similar narrow final range. GUIDE-NEI obtains the highest mean accuracy of 0.851271, followed by CGP-UCB at 0.851262. GUIDE-UCB obtains 0.851178, compared with

0.850973 for independent UCB, while independent NEI obtains 0.851145. The trajectories fluctuate more during the earlier rounds and become increasingly concentrated later.

Across the three tasks, GUIDE-NEI achieves the best average rank of 1.67, while GUIDE-UCB ranks third at 2.67. Both variants improve upon their matched independent baselines on all three tasks. GUIDE-TS has an average rank of 9.67, so the real-world evidence is strongest for the UCB and NEI instantiations.

The absolute gaps among the strongest methods are small on Activity Recognition and FedHPO, so the results indicate comparable performance among the leading methods rather than clear superiority over every collaborative baseline.

E Ablation Study Details

We conduct the ablation study using GUIDE-UCB on the same four representative benchmarks, namely Ackley, Rastrigin, Zakharov, and Michalewicz. The main study uses Level 2 heterogeneity, while the same configurations are additionally evaluated under Level 3 to examine whether their effects persist as cross-agent transferability decreases. All configurations use the same local GP models, posterior sampling procedure, initialization, evaluation budget, UCB decision rule, and repetition seeds. The Level 2 final results are reported in Table 3 in the main text.

Ablation variants. "Vanilla UCB" removes the complete federated distributional-exchange and uncertainty-intervention pipeline and performs standard local UCB optimization independently for each agent. It serves as a reference for measuring the overall benefit of federated guidance.

"Single-Gaussian belief" replaces the DPGMM-based representation with a single diagonal Gaussian fitted directly to all posterior-optimum samples, while retaining the same posterior sampling procedure and uploaded message format. This variant removes the multimodal representation of the distribution over optimum locations while preserving the remaining GUIDE pipeline.

"w/o value-aware reweighting" removes the standardized conservative value score from the server weights, such that the merged components are weighted only by their normalized mixture masses.

"w/o component merging" skips server-side clustering and moment matching while retaining value-aware reweighting and subsequent component sampling. Thus, the original uploaded components directly enter the downstream weighting and redistribution procedure.

"w/o agent-specific sampling" samples one global subset of at most P components and broadcasts the same subset to all agents. The packet size and weighted sampling distribution are kept identical to those of the full method, so this variant removes only the agent-specific stochastic redistribution.

"Uniform uncertainty scaling" retains the complete distributional-exchange pipeline but removes the spatial localization of FI-GP. Specifically, the spatially varying guidance field $G_{n,t}(\mathbf{x})$ is replaced at the intervention stage by its domain-averaged value

$$\bar{G}_{n,t} = \frac{1}{Q} \sum_{q=1}^{Q} G_{n,t}(\mathbf{z}_q),$$

where $\{\mathbf{z}_q\}_{q=1}^{Q}$ is a fixed quasi-uniform set over $\mathcal{X}$. The corresponding uncertainty scaling becomes

$$S_{n,t}^{\mathrm{uni}}(\mathbf{x}) = 1 + \lambda_t \bar{G}_{n,t}.$$

This preserves the average magnitude of the uncertainty intervention while removing its spatial variation, thereby separating spatially targeted guidance from a uniform increase in exploration.

E.1 Level 2 Ablation Results

Table 3 and Figure 8 reveal a clear difference between spatial uncertainty intervention and a generic increase in posterior uncertainty. Full GUIDE-UCB achieves the best average rank of 1.50. Removing agent-specific sampling gives an average rank of 2.50, followed by removing value-aware reweighting at 2.75, removing component merging at 3.75, and replacing the DPGMM with a single

Gaussian at 4.50. Uniform uncertainty scaling and Vanilla UCB give the two weakest aggregate results, with average ranks of 6.25 and 6.75.

The strongest effect comes from spatial localization. Full GUIDE-UCB, Uniform uncertainty scaling, and Vanilla UCB obtain final regrets of 8.4785, 19.1366, and 19.4350 on Ackley, and 58.3082, 80.3952, and 80.4771 on Rastrigin. On Zakharov, the corresponding values are 22.1717, 52.0749, and 49.8750, while on Michalewicz they are 5.8629, 6.1078, and 6.1248. The trajectories show the same pattern. Uniform uncertainty scaling stays close to Vanilla UCB on Ackley and Rastrigin and loses most of the early convergence advantage of Full GUIDE-UCB.

This ablation isolates an important part of FI-GP. The uniform variant keeps the complete distributional-exchange pipeline and preserves the average magnitude of the intervention, but removes its spatial variation. Its weak performance shows that a uniform increase in uncertainty is not sufficient. The main benefit comes from increasing uncertainty selectively in regions supported by the transferred distributions. FI-GP increases exploration around regions supported by the received distributions while leaving other parts of the search space much less affected.

The remaining ablations have smaller and more task-dependent effects. Replacing the DPGMM with a single Gaussian gives an average rank of 4.50. A single Gaussian can represent one broad promising region but cannot retain multiple separated modes when the posterior samples support several possible optimum locations. Component merging removes redundant components so that multiple entries in the limited downlink packet do not repeatedly represent the same region. Removing this step increases the average rank to 3.75.

Agent-specific sampling and value-aware reweighting affect the result less uniformly under Level 2. Removing agent-specific sampling gives the second-best average rank of 2.50 and even obtains 21.5431 on Zakharov, compared with 22.1717 for the full method. Removing value-aware reweighting obtains 8.4753 on Ackley, compared with 8.4785 for Full GUIDE-UCB. These reversals show that neither mechanism must improve every individual function. Value-aware reweighting makes component selection more selective, while agent-specific sampling gives different agents different subsets of the global components.

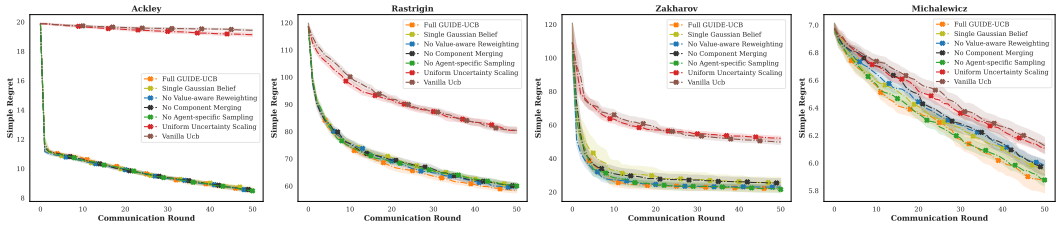

(a) Ackley (b) Rastrigin (c) Zakharov (d) Michalewicz

Figure 8: Convergence trajectories for the GUIDE-UCB ablation study under Level 2 heterogeneity. Curves show the mean average simple regret over 10 independent runs, and shaded regions denote ± 1 standard error of the mean across runs. Lower is better.

E.2 Level 3 Ablation Results

Under Level 3 heterogeneity, the ablation results become much more benchmark dependent. Full GUIDE-UCB nevertheless retains the best aggregate rank of 2.50. The Single-Gaussian variant ranks second at 3.00, followed by the variants without agent-specific sampling and component merging at 3.75 and 4.00. Uniform uncertainty scaling obtains 4.50, Vanilla UCB obtains 5.00, and removing value-aware reweighting gives the weakest aggregate rank of 5.25.

The best ablation variant now differs across functions. Removing value-aware reweighting is best on Ackley with 20.1612, while Full GUIDE-UCB obtains 20.1739. Uniform uncertainty scaling is best on Rastrigin at 90.2925, followed by the variant without agent-specific sampling at 90.7210 and Full GUIDE-UCB at 90.8849. Full GUIDE-UCB is best on Zakharov at 92.9444, whereas Vanilla UCB is best on Michalewicz at 7.3490. The Level 3 trajectories consequently overlap much more than their Level 2 counterparts.

The most notable change is the role of value-aware reweighting. Removing it gives an average rank of 2.75 at Level 2 but the worst rank of 5.25 at Level 3. Strong heterogeneity means that the server receives distributions from objectives that agree much less with one another. In this setting, mixture mass alone is not enough to distinguish a component that is merely frequent in posterior samples from one that is also supported by favorable local function values. The value-aware score provides this additional filter. Its effect is not uniform, as shown by the strong Ackley result without reweighting, but it becomes more important for avoiding weak aggregate performance across different tasks.

The Single-Gaussian variant ranks second overall, suggesting that the additional benefit of multi-modal representation is relatively modest in this Level 3 study. Uniform uncertainty scaling is best on Rastrigin, showing that spatial localization is not beneficial on every individual function. Under severe heterogeneity, filtering unreliable transferred components becomes more important because the received spatial information is less consistently relevant to each local objective.

Full GUIDE-UCB achieves the best Level 3 average rank of 2.50 despite not being the best variant on every benchmark. Its advantage is therefore primarily one of robustness across different objective landscapes.

Table 11: Ablation study of GUIDE-UCB on four representative synthetic benchmarks under Level 3 heterogeneity. Entries report the final mean average simple regret over 10 independent runs; lower is better. The last row reports the average rank across the four benchmarks. The best and second-best results are shown in bold and underlined, respectively.

Benchmark	Vanilla UCB	Single-Gaussian belief	w/o value-aware reweighting	w/o component merging	w/o agent-specific sampling	Uniform uncertainty scaling	Full GUIDE-UCB
Ackley	20.5073	20.2321	**20.1612**	20.2473	20.2122	20.4661	<u>20.1739</u>
Rastrigin	95.6377	92.1619	94.3347	92.5862	<u>90.7210</u>	**90.2925**	90.8849
Zakharov	101.9986	<u>95.9903</u>	109.8422	96.5643	98.2310	103.0493	**92.9444**
Michalewicz	**7.3490**	<u>7.3529</u>	7.5093	7.3660	7.4075	7.4026	7.3793
Average rank	5.00	<u>3.00</u>	5.25	4.00	3.75	4.50	**2.50**

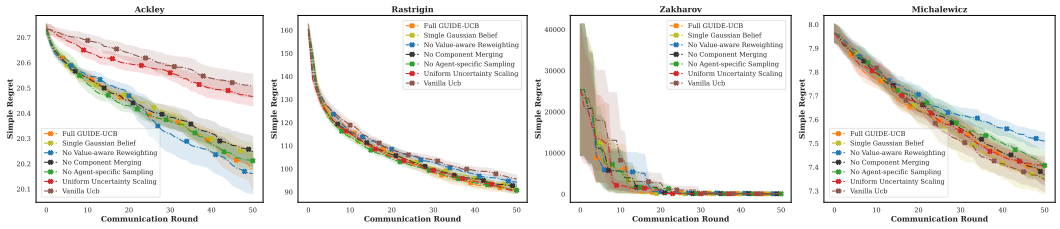

(a) Ackley (b) Rastrigin (c) Zakharov (d) Michalewicz

Figure 9: Convergence trajectories for the GUIDE-UCB ablation study under Level 3 heterogeneity. Curves show the mean average simple regret over 10 independent runs, and shaded regions denote ± 1 standard error of the mean across runs. Lower is better.

F Sensitivity Analysis

We evaluate the sensitivity of GUIDE-UCB under Level 3 heterogeneity on Rastrigin and Michalewicz. Level 3 has the lowest empirical task similarity ($\overline{\mathrm{Sim}} = 0.081$) and therefore provides a demanding setting for examining whether the main GUIDE hyperparameters remain well behaved when cross-agent agreement is weak. We vary one parameter at a time while keeping all remaining settings fixed.

The tested parameters are the number of posterior-optimum samples M, the downlink packet size P, the maximum intervention strength $\lambda_{\max}$, and the component merging threshold δ_{merge}. The complete parameter grids are reported in Table 12.

Table 13 reports the final average simple regret for each parameter setting. For each parameter group, only the parameter under study is varied, while all remaining hyperparameters are fixed at

Table 12: Hyperparameter grids used in the Level 3 sensitivity analysis. Bold entries denote the default settings.

Parameter	**Tested values**	**Default**
posterior-optimum samples M	$\{100, 250, \mathbf{500}, 750, 1000\}$	500
Downlink packet size P	$\{1, 3, \mathbf{5}, 8, 10\}$	5
Maximum guidance $\lambda_{\max}$	$\{0.25, 0.5, \mathbf{1.0}, 2.0\}$	1.0
Merging threshold δ_{merge}	$\{0.01, 0.02, 0.03, 0.04, \mathbf{0.05}\}$	0.05

Table 13: Sensitivity results of GUIDE-UCB on Rastrigin and Michalewicz under Level 3 heterogeneity. Entries report the final mean average simple regret over 10 independent runs; lower is better.

Parameter	**Value**	**Rastrigin**	**Michalewicz**
M	100	92.6392	7.3885
	250	91.9958	7.4102
	500	90.8849	7.3793
	750	91.3377	7.3489
	1000	93.1664	7.3473
P	1	92.5324	7.3935
	3	91.4631	7.3586
	5	90.8849	7.3793
	8	93.0066	7.3982
	10	93.2827	7.3982
$\lambda_{\max}$	0.25	90.7113	7.3563
	0.5	92.1546	7.3841
	1.0	90.8849	7.3793
	2.0	95.0430	7.4711
δ_{merge}	0.01	89.8604	7.3847
	0.02	90.5706	7.3613
	0.03	90.4103	7.3823
	0.04	89.5435	7.4494
	0.05	90.8849	7.3793

their default values. Figures 10 and 11 show the corresponding convergence trajectories over the complete BO horizon.

The effect of M differs between the two benchmarks. On Rastrigin, $M = 500$ gives the lowest final regret of 90.8849, while increasing the sample size to 750 and 1000 gives 91.3377 and 93.1664, respectively. Michalewicz shows the opposite tendency, with final regrets of 7.3489 and 7.3473 at $M = 750$ and $M = 1000$, compared with 7.3793 at $M = 500$. Increasing M beyond 500 therefore does not provide a consistent benefit across the two benchmarks. A moderate number of posterior-optimum samples is sufficient for Rastrigin, while Michalewicz benefits slightly from a finer Monte Carlo approximation of the distribution over optimum locations.

The packet size P also shows a task-dependent optimum. Rastrigin performs best at $P = 5$, with a final regret of 90.8849, whereas Michalewicz performs best at $P = 3$, with 7.3586. Increasing the packet size further degrades both benchmarks. At $P = 8$ and $P = 10$, Rastrigin reaches 93.0066 and 93.2827, while Michalewicz reaches 7.3982 for both settings. A larger packet gives each agent access to more global components, but under severe heterogeneity some of these additional components may correspond to regions that are poorly aligned with the local objective. Increasing P therefore does not necessarily improve the quality of the guidance. Moderate packet sizes also keep the communication cost bounded while limiting the amount of potentially conflicting global information received by each agent.

The clearest common pattern appears for $\lambda_{\max}$. Among the tested values, $\lambda_{\max} = 0.25$ gives the lowest final regret on both benchmarks, with 90.7113 on Rastrigin and 7.3563 on Michalewicz. In

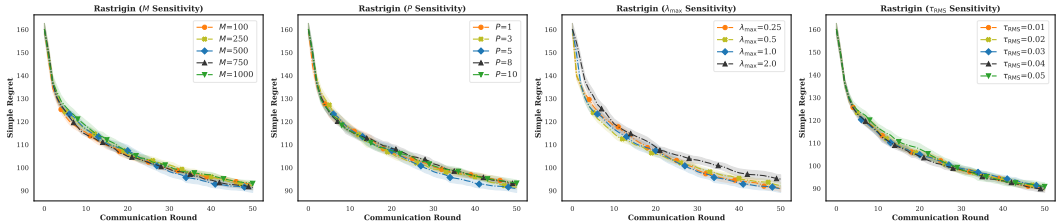

(a) Posterior-optimum samples M (b) Packet size P (c) Guidance strength $\lambda_{\max}$ (d) Merging threshold δ_{merge}

Figure 10: Sensitivity of GUIDE-UCB to its main hyperparameters on Rastrigin under Level 3 heterogeneity. Curves show the mean average simple regret over 10 independent runs, and shaded regions denote ± 1 standard error of the mean across runs. Lower is better.

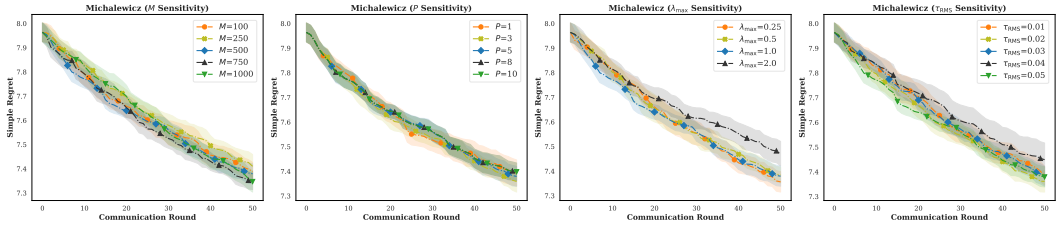

(a) Posterior-optimum samples M (b) Packet size P (c) Guidance strength $\lambda_{\max}$ (d) Merging threshold δ_{merge}

Figure 11: Sensitivity of GUIDE-UCB to its main hyperparameters on Michalewicz under Level 3 heterogeneity. Curves show the mean average simple regret over 10 independent runs, and shaded regions denote ± 1 standard error of the mean across runs. Lower is better.

contrast, the strongest intervention, $\lambda_{\max} = 2.0$, gives the worst result on both benchmarks, reaching 95.0430 and 7.4711, respectively. The results show that strong federated guidance is undesirable when cross-agent agreement is weak. Increasing posterior uncertainty too aggressively gives transferred information too much influence on the local decision. A weaker intervention is sufficient to exploit useful global information while leaving more weight to the local posterior.

The merging threshold δ_{merge} has no common optimum across the two benchmarks. Rastrigin achieves its lowest final regret of 89.5435 at $\delta_{\text{merge}} = 0.04$, whereas Michalewicz achieves its lowest value of 7.3613 at $\delta_{\text{merge}} = 0.02$. This difference is expected because the threshold controls the spatial resolution of the global representation. A larger value merges more nearby components and produces a coarser representation of the promising regions, while a smaller value retains finer spatial distinctions. The preferred resolution therefore depends on the geometry of the objective.

Overall, the sensitivity study does not favor uniformly larger parameter settings. Increasing M, P, or $\lambda_{\max}$ does not consistently reduce final regret, and the preferred merging threshold is task dependent. The most consistent finding is that excessively strong uncertainty intervention should be avoided under severe heterogeneity.

G Communication Cost Analysis

We measure communication by the number of scalar values transmitted per agent per BO round. Independent TS, UCB, and NEI do not communicate. The following accounting uses the default synthetic setting $d = 10$, $D_{\text{RFF}} = 500$, diagonal GUIDE covariance, and $P = 5$.

FTS and FTS-DE. FTS uploads and downloads a D_{RFF}-dimensional random feature weight representation, giving 500 scalars in each direction. Under the four-region distributed-exploration implementation of FTS-DE, the corresponding cost is $4D_{\text{RFF}} = 2000$ scalars in each direction.

FMTBO. FMTBO uploads two GP hyperparameters and ten ranking-related values, and receives two updated hyperparameters. Its uplink and downlink costs are therefore 12 and 2 scalars, respectively.

CGP family. Each CGP agent uploads one d-dimensional design point, one LCB value, and one posterior-mean value, for a total of $d + 2 = 12$ scalars. Each agent receives 20 design points, giving $20d = 200$ downlink scalars.

GUIDE-FBO. With a diagonal covariance, one uploaded component contains a d-dimensional mean, a d-dimensional covariance vector, one mixture weight, and one standardized value score. Hence,

$$C_{\mathrm{up}}^{\mathrm{GUIDE}} = 2d + 2 = 22.$$

Each downlink component contains a mean, diagonal covariance, and global guidance weight. Therefore,

$$C_{\mathrm{down}}^{\mathrm{GUIDE}} = P_t(2d + 1) \leq P(2d + 1) = 5 \times 21 = 105.$$

Thus, the total communication cost is at most 127 scalars per agent per round under the default setting. For a full covariance implementation, the covariance contribution becomes $d(d + 1)/2$ scalars per component when only the symmetric entries are transmitted. The resulting communication costs under the default diagonal-covariance implementation are summarized in Table 4.

H Computational Efficiency Analysis

We report the mean wall-clock time per independent run for each method using the same software environment and hardware. Experiments are executed on an AMD Ryzen 9 7945HX platform with 32 logical processors and 16 physical cores. Reported times are averaged over the repeated runs and measured in seconds.

Table 14: Mean wall-clock time per independent run on the synthetic benchmarks under Level 2 heterogeneity, averaged over 10 runs. Unit: seconds.

Benchmark	TS	UCB	NEI	FTS	FTS-DE	FMTBO	CGP-TS	CGP-UCB	CGP-NEI	GUIDE-TS	GUIDE-UCB	GUIDE-NEI
Ackley	129.51	143.48	269.34	141.88	148.32	160.56	252.10	411.90	1811.38	267.19	275.87	402.82
Levy	111.27	125.76	262.50	133.12	124.19	256.59	237.84	263.90	765.01	215.37	217.87	499.63
Griewank	85.13	99.67	279.07	118.16	113.43	242.71	246.30	253.81	773.83	179.68	185.57	344.68
Rastrigin	120.42	163.72	290.74	128.78	121.39	297.36	265.75	359.93	1221.39	227.67	262.98	555.44
Weierstrass	126.20	140.59	250.92	141.79	149.75	171.29	246.86	347.38	1050.71	266.59	267.94	677.94
Ellipsoid	175.70	164.10	306.68	161.63	153.78	246.67	304.80	326.98	892.53	249.66	253.14	624.32
Sphere	167.19	174.37	327.94	178.33	164.62	266.78	257.60	275.84	814.01	223.04	234.56	534.97
Zakharov	115.02	113.23	224.83	115.94	120.83	225.49	242.18	240.12	449.03	233.55	240.50	562.81
Rosenbrock	65.79	78.15	172.91	78.58	89.00	154.95	154.68	163.35	333.75	194.66	201.80	439.28
Michalewicz	114.72	166.72	267.53	131.72	137.26	195.98	231.74	368.63	1406.76	264.11	294.16	742.14
Powell	82.17	90.78	186.99	97.34	107.10	167.66	185.97	185.55	296.17	227.31	234.32	477.27
Styblinski–Tang	83.22	88.35	216.58	101.46	106.42	183.35	224.13	274.83	1048.93	206.32	219.41	548.39

The synthetic objectives are inexpensive to evaluate, so the reported wall-clock times primarily reflect the computational overhead of the BO algorithms themselves. Averaged across the 12 Level 2 benchmarks, UCB, GUIDE-UCB, and CGP-UCB require approximately 129.1, 240.7, and 289.4 seconds per complete run, respectively. GUIDE-UCB is faster than CGP-UCB on 9 of the 12 benchmarks.

For the NEI family, GUIDE-NEI and CGP-NEI require approximately 534.1 and 905.3 seconds on average, respectively, with GUIDE-NEI being faster on 9 of the 12 benchmarks. GUIDE therefore introduces clear additional computation relative to independent local BO, mainly from posterior-optimum sampling, distribution fitting, and server-side processing, while remaining less expensive than the corresponding CGP variants on most of the tested benchmarks.

This additional computation is a practical cost when objective evaluations themselves are inexpensive. In expensive black-box applications, however, simulation, model training, or physical experimentation may dominate the total optimization time, in which case optimizer-side computation can represent a smaller fraction of the end-to-end cost. The relative runtime impact therefore depends on the expense of the application being optimized.

I ORTHOGONAL RANDOM FEATURES FOR POSTERIOR-OPTIMUM SAMPLING

GUIDE uses random Fourier features (Rahimi & Recht, 2007) with an orthogonal-block construction inspired by Yu et al. (2016) to generate approximate posterior paths. The local GP provides the posterior mean, while a Bayesian linear feature model provides the stochastic residual.

Feature representation. For a fixed agent and round, let $\boldsymbol{\ell}$ and σ_f^2 denote the fitted local GP's ARD lengthscales and output variance, and define $\boldsymbol{\Lambda} = \mathrm{diag}(\boldsymbol{\ell})$. We suppress the agent and round indices on the feature map for readability. Direction vectors $\mathbf{q}_j$ are obtained from rows of orthogonal matrices generated by QR decomposition of independent standard Gaussian matrices. For each direction, draw independent $u_j \sim \chi^2_d$, $v_j \sim \chi^2_5$, and $b_j \sim \mathrm{Unif}[0, 2\pi)$, and define

$$\boldsymbol{\omega}_j = \sqrt{u_j}\sqrt{\frac{5}{v_j}}\,\boldsymbol{\Lambda}^{-1}\mathbf{q}_j, \qquad \phi_j(\mathbf{x}) = \sqrt{\frac{2\sigma_f^2}{D_{\mathrm{ORF}}}}\cos(\boldsymbol{\omega}_j^\top \mathbf{x} + b_j).$$

The Student-t_5 radial scaling and ARD transformation follow the spectral parameterization of the local Matérn-$5/2$ kernel. The resulting feature map approximates its stationary covariance through $k(\mathbf{x}, \mathbf{x}') \approx \boldsymbol{\phi}(\mathbf{x})^\top \boldsymbol{\phi}(\mathbf{x}')$.

Posterior paths. We sample feature weights from their Gaussian posterior under the Bayesian linear feature model. Let $\widetilde{\mathbf{w}}_{n,t}^{(m)}$ denote such a draw conditioned on $\mathcal{D}_{n,t-1}$, and let $\boldsymbol{\mu}_{w,n,t-1}$ denote the corresponding posterior weight mean. To retain the fitted local GP posterior mean, we construct

$$\widehat{f}_{n,t}^{(m)}(\mathbf{x}) = \mu_{n,t-1}(\mathbf{x}) + \boldsymbol{\phi}(\mathbf{x})^\top \left(\widetilde{\mathbf{w}}_{n,t}^{(m)} - \boldsymbol{\mu}_{w,n,t-1}\right).$$

Conditional on the fitted model and feature map, the residual has zero mean. Thus, the paths preserve the local GP posterior mean while approximating its posterior covariance through the feature model.

Optimum distribution extraction and TS decisions. For optimum distribution extraction, we generate M paths with independent posterior weight draws, sharing a feature map and a uniformly sampled candidate set $\mathcal{X}_{\mathrm{cand}} \subset \mathcal{X}$. We obtain optimum-location samples as

$$\widehat{\mathbf{x}}_{n,t}^{(m)} \in \arg\max_{\mathbf{x}\in\mathcal{X}_{\mathrm{cand}}} \widehat{f}_{n,t}^{(m)}(\mathbf{x}),$$

and use these samples to fit the local DPGMM.

For GUIDE-TS, a fresh posterior weight draw defines a path whose centered residual is scaled by the guidance multiplier:

$$\widetilde{f}_{n,t}(\mathbf{x}) = \mu_{n,t-1}(\mathbf{x}) + S_{n,t}(\mathbf{x})\left(\widehat{f}_{n,t}(\mathbf{x}) - \mu_{n,t-1}(\mathbf{x})\right).$$

This path is optimized over $\mathcal{X}$ using continuous acquisition optimization. For GUIDE-UCB and GUIDE-NEI, the feature sampler is used for optimum distribution extraction, while acquisition evaluation uses the intervened local GP posterior. All features and sampled paths remain on the agent; only the selected distributional component is uploaded.