Technometrics

ISSN: (Print) (Online) Journal homepage: www.tandfonline.com/journals/utch20

Multi-Agent Collaborative Bayesian Optimization via Constrained Gaussian Processes

Qiyuan Chen, Liangkui Jiang, Hantang Qin & Raed Al Kontar

To cite this article: Qiyuan Chen, Liangkui Jiang, Hantang Qin & Raed Al Kontar (2025) Multi-Agent Collaborative Bayesian Optimization via Constrained Gaussian Processes, Technometrics, 67:1, 32-45, DOI: 10.1080/00401706.2024.2365732

To link to this article: https://doi.org/10.1080/00401706.2024.2365732

TECHNOMETRICS
2025, VOL. 67, NO. 1, 32–45
https://doi.org/10.1080/00401706.2024.2365732

Taylor & Francis
Taylor & Francis Group

Multi-Agent Collaborative Bayesian Optimization via Constrained Gaussian Processes

Qiyuan Chen[a], Liangkui Jiang[b], Hantang Qin[b], and Raed Al Kontar[a]

[a]University of Michigan, Ann Arbor, MI; [b]University of Wisconsin, Madison, WI

ABSTRACT

The increase in the computational power of edge devices has opened a new paradigm for collaborative analytics whereby agents borrow strength from each other to improve their learning capabilities. This work focuses on collaborative Bayesian optimization (BO), in which agents work together to efficiently optimize black-box functions without the need for sensitive data exchange. Our idea revolves around introducing a class of constrained Gaussian process surrogates, enabling agents to borrow informative designs from high-performing collaborators to enhance and expedite their optimization process. Our approach presents the first general-purpose collaborative BO framework that is compatible with any Gaussian process kernel and most of the known acquisition functions. Despite the simplicity of our approach, we demonstrate that it offers elegant theoretical guarantees and significantly outperforms state-of-the-art methods, especially when agents have heterogeneous black-box functions. Through various simulations and a real-life experiment in additive manufacturing, we showcase the advantageous properties of our approach and the benefits derived from collaboration.

ARTICLE HISTORY
Received December 2023
Accepted May 2024

KEYWORDS
Collaborative Bayesian optimization; Constrained Gaussian processes; Federated learning

1. Introduction

Bayesian optimization (BO) is a long-established yet prominent method for optimizing black-box functions. Abstractly, numerous real-world problems can be framed as finding the best outcome by manipulating a set of variables called *designs*. Often, the relationship between the outcome and designs is unknown, and hence, obtaining optimal designs can only be achieved via trial and error. For instance, clinical trials often manipulate medication dosages to uncover the most effective and safe treatment options. Likewise, manufacturers tweak production process parameters to minimize product defects. This involves selecting and testing a sequence of designs and observing their occasionally noisy outcomes, referred to as *responses* in the quest for the best design. Since these experiments are often expensive and the search regions are continuous, an exhaustive search is rarely plausible. BO is a family of methods within sequential optimal design (Friedman and Savage 1947) that wisely allocates a limited number of experiments through a *sequential process*. Its objective is to identify the best response with the fewest trials.

A typical BO algorithm first constructs a posterior belief of the underlying black-box function using a surrogate, then determines the next design to observe (i.e., do a trial) by optimizing a user-specified utility function, known as the acquisition function (AF), that quantifies the benefits of experimenting at a new design. Upon observing a response from the selected design, BO then updates its posterior belief and selects a new design. The process is repeated until the experimentation budget is exhausted or some exit condition is met. Needless to say, BO has found success in a wide variety of science and engineering disciplines, including but not limited to chemistry (Shields et al. 2021), physics (Duris et al. 2020), biology (Gonzalez et al. 2015), material science (Attia et al. 2020) and mechanical design (Gongora et al. 2020), amongst many others.

Traditionally, BO has assumed that all experimental designs are sequentially determined and observed by a single agent. However, in practice, resources (e.g., funding, materials, subjects, computational power) for conducting experiments are often distributed among different agents. Instead of each agent conducting their sequential experiments in isolation, collaboration between agents can potentially expedite the process and help individuals find better designs within limited resource budgets. For instance, consider a group of 3D-printing workshops fabricating a material. The printing time of this material depends on a set of process parameters specified by users, which is a typical BO scenario. The goal of collaboration is to enable all the workshops to borrow strength from each other to fast-track and improve the optimal design process of their process parameters.

Unfortunately, collaboration comes with its unique challenges. First, collaborators could be heterogeneous, meaning that knowledge borrowed from others might not always be helpful, or, in some cases may even be misleading. In the previous example, optimal process parameters may differ among workshops due to varying operational environments. Second, in many real-world settings, such as healthcare or security, the direct sharing of raw data is restricted due to privacy concerns, communication burdens, intellectual property protection, or regulations (Kontar et al. 2021).

CONTACT Raed Al Kontar alkontar@umich.edu University of Michigan, Ann Arbor 48103, MI.
Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/TECH.

Contributions:. This work aims to address the exact challenge above. Our contribution can be summarized as follows:

(a) *Algorithm:* We propose the first general-purpose collaborative BO framework that is compatible with most existing AFs and is provably robust to heterogeneity across collaborative agents. At a high level, our approach hinges upon defining a constrained Gaussian process ($\mathcal{CGP}$) where an agent exploits designs from strongly performing collaborators to tweak their surrogates toward potential regions of improvement.
(b) *Compatibility:* We present an easy-to-implement approach to applying our method under three well-known categories of AFs. Our approach also works for synchronous and asynchronous settings, with or without a central server.
(c) *Theory:* We prove, regardless of heterogeneity, that the regret of our approach is sub-linear under a specific class of AFs.
(d) *Practical performance:* Empirically, we demonstrate that our approach can significantly outperform the state-of-the-art in optimality gap and instantaneous regret. As a proof of concept, we apply our algorithm to a 3D printing process parameter selection problem to highlight the value of collaboration.

2. Related Work

BO is a branch of sequential design problems (Gramacy 2020) that aims to optimize black-box functions. The methodological philosophy of BO has a rich history. Hotelling (1941) first proposed using a surrogate model to estimate the maximum of a black-box function. However, their method picks all the design points to experiment on in one shot, which was considered inefficient by Friedman and Savage (1947) as they argued that experiments should be more focused on exploring near-optimal regions. As such, Friedman and Savage (1947) introduced a *sequential design* approach instead, where the new design points were determined iteratively using a cyclic coordinate descent heuristic. Yet, the coordinate descent heuristic fails to be a global optimization scheme since it can get stuck in local-optimal solutions. Various heuristics were proposed afterward until the Bayesian heuristic (Raiffa and Schlaifer 1961) took dominance in the field. The Bayesian heuristic picks experiments by maximizing a user-specified utility function, which was later referred to as the acquisition function (AF). Finally, with all the aforementioned building blocks, the BO framework was formalized by Kushner (1963), where they used a Gaussian Process ($\mathcal{GP}$) surrogate and the probability of improvement (PI) AF.

Since then, the modern BO paradigm has used a surrogate model to characterize the posterior belief of a black-box function and then chooses designs sequentially by maximizing an AF. The surrogate model has been predominantly a $\mathcal{GP}$, but the choice of AF has a wide variety and remains one of the most active and impactful research fields in BO. Famous AF's include Thompson sampling (TS; Thompson 1933), expected improvement (EI; Jones, Schonlau, and Welch 1998), probability of improvement (PI; Kushner 1963), upper confidence bound (UCB; Srinivas et al. 2009), knowledge gradient (KG; Wu and Frazier 2016), Noisy Expected Improvement (NEI; Letham et al. 2019), max-value entropy search (MES; Wang and Jegelka 2017), etc. Besides research on designing new and improved AFs (Wu and Frazier 2016; Wang and Jegelka 2017; Letham et al. 2019), recent advances in BO include scaling BO methods to high-dimensional input via subspace projections (Moriconi, Deisenroth, and Sesh Kumar 2020), extending BO to multi-objective (Daulton et al. 2022) and multi-fidelity settings (Irshad, Karsch, and Döpp 2021), amongst others.

Despite many advances, BO was only recently extended to a collaborative setting under the notion of federated BO (FBO; Dai, Low, and Jaillet 2020, 2021). This literature was inspired by the surge of computational and data-collecting power at the edge, which has set forth a new paradigm of collaborative yet privacy-preserving learning coined as federated learning (Kontar et al. 2021). Indeed, FBO shares commonalities with earlier literature on batch BO (Azimi, Fern, and Fern 2010; Duan et al. 2017; Hunt 2020). However, batch BO focuses on optimizing a single objective using one surrogate learned from a centralized dataset. In contrast, in FBO, data comes from potentially heterogeneous agents, and each aims to optimize their own function; as such, there may be no surrogate that fits all. The other related but different line of literature is on Transfer BO (Bai et al. 2023) where data from relevant tasks is used to accelerate current tasks. Although Transfer BO method deals with heterogeneity in the system, they assume data of every agent is available to every other agent in the system. Unlike batch BO or transfer BO, in our setting, data is generated locally across multiple agents and is unavailable (nor should it be shared) centrally.

To the best of our knowledge, the first FBO approach (Dai, Low, and Jaillet 2020) proposed federated Thompson sampling (FTS), which was later extended to FTS with distributed exploration (FTSDE; Dai, Low, and Jaillet 2021). The key idea is to share function realizations of the posterior belief using random Fourier features and then dwindle collaboration as more data is collected. More importantly, this work highlighted the ability of agents to benefit from collaboration without sharing raw data.

One major limitation of existing literature in FBO is that they restrict the choice of AF to only TS, leaving along the rich literature in AFs. The family of AFs defines their searching heuristics, which mainly involve balancing exploration and exploitation. Different strategies yield different performances across different tasks. It is also common for BO studies to design AFs for specific practical tasks. In short, there is no single AF that fits all, but one needs to pick the right AF for their own tasks. To this end, this work embraces the variety in AFs by envisioning a new perspective for collaboration that enables agents to tweak their surrogates toward potential regions of improvement through a constrained Gaussian Process.

3. Setup and Notation

Consider a group of N agents indexed by n where each agent aims to optimize their black-box function $f_n : \mathcal{X} \subset \mathbb{R}^D \to \mathbb{R}$. In many applications, the functions $f_n, \forall n \in [N]$ share commonalities across agents but may feature some idiosyncrasy as agents operate possibly under different conditions or locations. Without loss of generality, we consider a maximization problem throughout this article. For each agent n, the goal is to find an optimal design that maximizes their own black-box function, that is, $\boldsymbol{x}_n^* \in \arg\max_{\boldsymbol{x}\in\mathcal{X}} f_n(\boldsymbol{x})$. BO is a family of algorithms

that solves this problem by sequential experimentation. Again, without the loss of generality, we assume the experiments are performed one by one iteratively to simplify notations. Therefore, at every round $t \in [T]$, each agent n decides on the next design $\boldsymbol{x}_{n,t+1}$ to sample and then runs an experiment to observe the response value $f_n(\boldsymbol{x}_{n,t+1})$. Realistically, agents can only obtain a noisy version of the true function, $y_n(\boldsymbol{x}_{n,t+1}) = f_n(\boldsymbol{x}_{n,t+1}) + \varepsilon_{n,t+1}$, where $\varepsilon_{n,t+1}$ denotes observational noise.

To this end, BO selects designs at each round t by optimizing an AF $\alpha(\cdot;\cdot)$. Specifically in our problem, agents select their own next design by optimizing their own AF. The AF takes two inputs, namely a design and a posterior belief. It quantifies the benefits one would gain (according to one's posterior belief) if an experiment is conducted at the given design point. To characterize the posterior belief, BO resorts to building a surrogate model $F_{n,t}$ at round t that estimates the function and quantifies the uncertainty in f_n based on a currently available dataset $\mathcal{D}_{n,t} := \{(\boldsymbol{x}_{n,1}, y_{n,1}), \ldots, (\boldsymbol{x}_{n,t}, y_{n,t})\}$ at time t.

With that, BO selects $\boldsymbol{x}_{n,t}$ to be $\arg\max_{x\in\mathcal{X}} \alpha(x; F_{n,t})$. Predominantly, since Kushner (1963), $F_{n,t}$ has been modeled as a posterior belief characterized by a $\mathcal{GP}$, due to a $\mathcal{GP}$'s elegant formulation and natural uncertainty quantification. A $\mathcal{GP}$ is defined by a positive-definite kernel $k_n(\cdot,\cdot)$ that imposes a prior belief on the black-box function, which is later conditioned on the observed data $\mathcal{D}_{n,t}$ to obtain the posterior belief. Hereon, and for ease of later development, we denote the posterior belief $F_{n,t} \triangleq \mathcal{GP}(\mathcal{D}_{n,t})$. Now, to set the notation, for agent n at time t, define $k_{n,t}(x) = [k_n(x, \boldsymbol{x}_{n,1}), \cdots, k_n(x, \boldsymbol{x}_{n,t})]^\intercal$ and the Gram matrix $K_{n,t} = [k_{n,t}(\boldsymbol{x}_{n,1}), \ldots, k_{n,t}(\boldsymbol{x}_{n,t})]$. By a $\mathcal{GP}(\mathcal{D}_{n,t})$ construction, we have that:

$$\begin{bmatrix} y_{n,:t} \\ f_n(x) \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} K_{n,t} + \hat{\lambda}_t I & k_{n,t}(x) \\ k_{n,t}(x)^\intercal & k_{n,t}(x,x) \end{bmatrix}\right),$$

where $y_{n,:t} = [y_{n,1}, \ldots, y_{n,t}]^\intercal$ and $\hat{\lambda}_t$ is a nugget effect that provides regularization and alleviates numerical problems for $\mathcal{GP}$s (Matheron 1963). By marginalizing all the possible function realizations, we get the agent's posterior belief $f_n(x)|\mathcal{D}_{n,t} \sim \mathcal{N}(\mu_{n,t}(x), \sigma^2_{n,t}(x))$, where

$$\mu_{n,t}(x) = k_{n,t}(x)^\intercal (K_{n,t} + \hat{\lambda}_t I)^{-1} y_{n,:t},$$
$$\sigma^2_{n,t}(x) = k_n(x,x) - k_{n,t}(x)^\intercal (K_{n,t} + \hat{\lambda}_t I)^{-1} k_{n,t}(x).$$

4. Methodology

As discussed above, an agent's choice of the design point to sample is based on their posterior belief in f_n. Yet, here lies the key opportunity enabled by collaboration: Can agents share information that allows them to improve and fast-track their optimal design process by guiding their posterior beliefs toward better designs? Importantly, the answer to this question should meet two critical constraints. The first is designing for heterogeneity. Even if raw observations were shared, agents could not simply update their surrogates by augmenting their local datasets with data from other agents, as the underlying functions may be heterogeneous across agents. The second constraint is privacy. Privacy concerns may prohibit the direct sharing of raw observations among agents. Indeed, for agents to participate in a collaborative process, it may be imperative to preserve their raw data.

To address the aforementioned challenges, we propose a natural idea to share some key statistics $E_{n,t}$ in lieu, where $E_{n,t}$ is extra information received by agent n from other agents $n' \in [N \setminus n]$ at round t. The key challenge lies in defining $E_{n,t}$ and determining how it should be effectively employed. For instance, when $F_{n,t}$ is specified as a $\mathcal{GP}$, conditioning on anything other than data points breaks the closed-form posterior of the $\mathcal{GP}$. The following sections aim to address two critical questions: (i) What should be included in $E_{n,t}$? (ii) Given $E_{n,t}$, how can we update the posterior belief $F_{n,t}|E_{n,t}$ and, consequently, the AF?

4.1. Collaboration by Sharing Constraints

As mentioned in the literature review, one fundamental design philosophy of BO is to allocate designs to near-optimal regions (Friedman and Savage 1947). Since we believe the black-box functions may have commonalities across agents, they may share similar near-optimal regions. Naturally, agents can distributively explore the search space and collectively locate the near-optimal region. As soon as agents locate the near-optimal regions, the problem is largely simplified to finding a personalized solution for each agent within the region.

Following this rationale, we design $E_{n,t}$ to entail information about the near-optimal regions. We let $E_{n,t}$ be a set of *synthetic* designs shared from other agents $n' \in [N \setminus n]$ to agent n, such that every borrowed design $\boldsymbol{x}^+_{n',t} \in E_{n,t}$ satisfies the constraint $f_n(\boldsymbol{x}^+_{n',t}) > \kappa_{n,t}$ for a constant $\kappa_{n,t}$ picked by agent n. Therefore, any design in $E_{n,t}$ points out to a region that has a value of at least $\kappa_{n,t}$. In this article, we let $\kappa_{n,t}$ be the maximum posterior mean of agent n (see Figure 1). That is

$$\kappa_{n,t} := \max_x \mu_{n,t}(x).$$

Therefore, every design $\boldsymbol{x}^+_{n',t} \in E_{n,t}$ points out a potential design with a better response compared to agent n's current best, that is, $\max_x \mu_{n,t}(x)$, according to their posterior belief.

Now, notice that even if all agents have the same $f = f_n \,\forall\, n$, the shared information $f(\boldsymbol{x}^+_{n',t}) > \kappa_{n,t}$ may still be incorrect as agent n' only has a posterior belief $F_{n',t}$ and not the true f. Since observations are noisy and the surrogate models are never exact, one can never guarantee that the shared information is correct, but we can foster the integrity of the information received by others by being conservative. Therefore, we restrict other agents n' to share a design $\boldsymbol{x}^+_{n',t}$ to $E_{n,t}$ only if their lower confidence bound (LCB) at $\boldsymbol{x}^+_{n',t}$ is greater than $\kappa_{n,t}$:

$$\mu_{n',t}(x) - \eta_t \cdot \sigma_{n',t}(x) > \kappa_{n,t}\,, \tag{1}$$

where η_t is a user-specified confidence level. In other words, agent n' believes $\mu_{n',t}(\boldsymbol{x}^+_{n',t}) \geq \kappa_{n,t}$ with high confidence.

To highlight this idea, a simple illustration is provided in Figure 1. Agent n' will share a design $\boldsymbol{x}^+_{n',t}$ from the shaded region such that its lower confidence bound is greater than $\kappa_{n,t}$ chosen by agent n. Clearly, there may be many, or none, $\boldsymbol{x}^+_{n',t}$ that satisfy (1).

In this article, we assume that each agent contributes at most one design in $E_{n,t}$. To do so, agent n' needs to find a design for agent n such that (1) is satisfied. Our choice for such design to share (when it exists) is the LCB maximizer $\boldsymbol{x}^+_{n',t}$ (see Figure 1):

$$\boldsymbol{x}^+_{n',t} \in \arg\max_{x\in\mathcal{X}} \left\{\mu_{n',t}(x) - \eta_t \cdot \sigma_{n',t}(x)\right\}$$

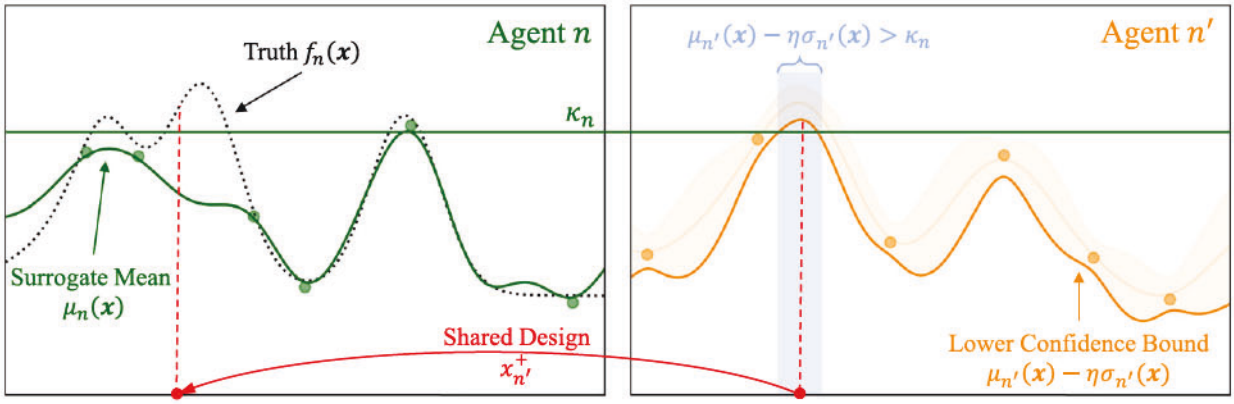

Figure 1. Agent n' chooses to share a design $x^+_{n'}$ from the shaded region where the lower confidence bound is greater than κ_n chosen by agent n (subscript t is dropped here).

Critically, two observations can be made for agent n''s LCB maximizer, $x^+_{n',t}$. First, if there exists any x satisfying inequality (1), then $x^+_{n',t}$ must also satisfy (1). Hence, to solve the feasibility problem posed by inequality (1), it suffices to find the maximum LCB $(\mu_{n',t}(x^+_{n',t}) - \eta_t \cdot \sigma_{n',t}(x^+_{n',t}))$ from agent n' and compare it with $\kappa_{n,t}$. Second, the maximum posterior mean of agent n' is promised to be greater than that of agent n (i.e., $\kappa_{n't} > \kappa_{n,t}$), meaning agent n' takes the lead in the optimization process, so we call agent n' a *leader* of agent n. Therefore, letting $\mathcal{A}_{n,t} \subset [N \setminus n]$ be the set of all the leaders of agent n at round t, the shared information $E_{n,t}$ is given by

$$E_{n,t} := \{x^+_{n',t} : n' \in \mathcal{A}_{n,t}\}.$$

Recall that for agent n to obtain $E_{n,t}$, the agent only needs to compare $\kappa_{n,t}$ with the maximum LCB from agents $n' \in [N \setminus n]$. If the maximum LCB from agent n' satisfies

$$\mu_{n',t}(x^+_{n',t}) - \eta_t \cdot \sigma_{n',t}(x^+_{n',t}) > \kappa_{n,t}\,, \tag{2}$$

then $x^+_{n',t}$ joins the set $E_{n,t}$ and is borrowed by agent n. This is a rather simple comparison that can be done privately in either a centralized or decentralized manner. We provide more details on this comparison in Section 4.4.3.

Hereon, after finding $E_{n,t}$, agent n presumes that for all $x^+_{n',t} \in E_{n,t}$, there holds $f_n(x^+_{n',t}) > \kappa_{n,t}$. Even with our conservative choice of $E_{n,t}$, this may still not hold because of unknown heterogeneity and surrogate estimation error. Nonetheless, as will become clear shortly, agent n will hedge against this heterogeneity through a rejection sampling scheme. Such a scheme will lead to provable convergence (and strong empirical performance) no matter how heterogeneous clients are.

4.2. Learning the AF with the Shared Constraints

A natural way to exploit the shared information $E_{n,t}$ is for agent n to condition their surrogate on it. In our case, because $E_{n,t}$ encompasses a set of constraints, the updated posterior belief

$$F^+_{n,t} \triangleq \mathcal{GP}(\mathcal{D}_{n,t})|E_{n,t}, \tag{3}$$

becomes a constrained Gaussian Process ($\mathcal{CGP}$). With the extra information $E_{n,t}$, we expect the updated posterior belief $\mathcal{GP}$ to better guide the AF in the search for optimal designs. Unfortunately, being a $\mathcal{CGP}$, the updated posterior renders the calculation of most existing BO AFs intractable. For instance, EI will lose its closed form after conditioning on $E_{n,t}$.

Nevertheless, although $\mathcal{CGP}$ itself does not have a closed-form posterior, we can still draw function realization from it and approximate the conditioned AFs using Monte Carlo methods. We demonstrate how this applies to three primary classes of AFs.

4.2.1. Acquisition Functions that Sample Function Realizations

We start by introducing how to learn the AF by drawing samples from $F^+_{n,t}$. Thompson Sampling (TS; Thompson 1933) is a typical AF in this class that selects design points by drawing sample(s) from the posterior belief of optimal designs. Formally, when the posterior belief is given as $F^+_{n,t}$, the AF for TS can be simply denoted as

$$\alpha(x; F^+_{n,t}) = g_{n,t,s}(x)\,, \tag{4}$$

where $\{g_{n,t,s}\}_{s\in[S]}$ are function realizations drawn from $F^+_{n,t}$. The subscript s here denotes the sample index within a batch of samples. Therefore, after drawing $g_{n,t,s}$, the next design to observe for agent n is simply the AF maximizer $\arg\max_{x\in\mathcal{X}} \alpha(x; F^+_{n,t}) = \arg\max_{x\in\mathcal{X}} g_{n,t,s}(x)$.

The challenge is drawing $g_{n,t,s} \sim F^+_{n,t}$. Letting $p(\cdot)$ be a probability density function, we have that $g_{n,t,s} \sim p(f_{n,t}|\mathcal{D}_{n,t}, E_{n,t})$, which is not trivial as the distribution is no longer Gaussian after conditioning on $E_{n,t}$. To this end, we introduce a set of ancillary variables:

$$\mathcal{D}^+_{n,t} := \{(x^+_{n',t}, f_n(x^+_{n',t})) : x^+_{n',t} \in E_{n,t}\}\,. \tag{5}$$

$\mathcal{D}^+_{n,t}$ denotes agent n's unknown ground truth data at $E_{n,t}$. This simple trick allows us to obtain samples $\{g_{n,t,s}\}_{s\in[S]}$ by sampling from the joint distribution $(f_{n,t}, \mathcal{D}^+_{n,t}) \sim p(f_{n,t}, \mathcal{D}^+_{n,t}|\mathcal{D}_{n,t}, E_{n,t})$ instead. More specifically, we have that:

$$\begin{aligned} p(f_{n,t}, \mathcal{D}^+_{n,t}|\mathcal{D}_{n,t}, E_{n,t}) &= p(f_{n,t}|\mathcal{D}_{n,t}, E_{n,t}, \mathcal{D}^+_{n,t}) \cdot p(\mathcal{D}^+_{n,t}|\mathcal{D}_{n,t}, E_{n,t}) \\ &= p(f_{n,t}|\mathcal{D}_{n,t}, \mathcal{D}^+_{n,t}) \cdot p(\mathcal{D}^+_{n,t}|\mathcal{D}_{n,t}, E_{n,t}) \\ &\propto \underbrace{p(f_{n,t}|\mathcal{D}_{n,t}, \mathcal{D}^+_{n,t})}_{\text{Fantasy Model}} \cdot \underbrace{p(E_{n,t}|\mathcal{D}^+_{n,t})}_{\text{Rejection}} \cdot \underbrace{p(\mathcal{D}^+_{n,t}|\mathcal{D}_{n,t})}_{\text{Fantasy Samples}}\,, \end{aligned}$$

where the second equality is because $\mathcal{D}^+_{n,t}$ is a sufficient statistic for $E_{n,t}$, and the last line is by Bayes' theorem. This decomposition outlines that the sampling of a function realization $g_{n,t,s}$ can be realized in three steps using rejection sampling.

Namely, in the first step, we can take $S_{\texttt{raw}}$ samples from the posterior $p(\mathcal{D}^+_{n,t}|\mathcal{D}_{n,t})$ given by the model $\mathcal{GP}(\mathcal{D}_{n,t})$. Within each sample s, we denote the sampled response at $\boldsymbol{x}^+_{n',t}$ as $g_{n,t,s}(\boldsymbol{x}^+_{n',t})$. Since we only need responses at $g_{n,t,s}(\boldsymbol{x}^+_{n',t})$, the samples can be easily taken from the $\mathcal{GP}$ predictive distribution. As a second step, and since an agent presumes all $\boldsymbol{x} \in E_{n,t}$ satisfies $f_n(\boldsymbol{x}) > \kappa_{n,t}$, we retain $g_{n,t,s}$ only if

$$g_{n,t,s}(\boldsymbol{x}^+_{n',t}) > \kappa_{n,t}, \quad \forall \boldsymbol{x}^+_{n',t} \in E_{n,t}. \tag{6}$$

We note that any accepted $g_{n,t,s}$ should satisfy *all* constraints posed by $E_{n,t}$. As such, for each (accepted) sample index s, we now have a new dataset consisting of the retained samples

$$\begin{aligned} &\mathcal{S}^+_{n,t,s} \sim p(\mathcal{D}^+_{n,t}|\mathcal{D}_{n,t}, E_{n,t}), \\ &\text{where } \mathcal{S}^+_{n,t,s} \triangleq \{(\boldsymbol{x}^+_{n',t}, g_{n,t,s}(\boldsymbol{x}^+_{n',t})) : \boldsymbol{x}^+_{n',t} \in E_{n,t}\}. \end{aligned} \tag{7}$$

The retained samples $\mathcal{S}^+_{n,t,s}$ reflect the "imagination" of the unknown truth $\mathcal{D}^+_{n,t}$ given all the information agent n has (i.e., $\mathcal{D}_{n,t}$ and $E_{n,t}$). Hence, we call such a retained dataset $\mathcal{S}^+_{n,t,s}$ as a *fantasy dataset*. For a fantasy dataset $\mathcal{S}^+_{n,t,s}$, one can construct a new $\mathcal{GP}$

$$M_{n,t,s} := \mathcal{GP}(\mathcal{D}_{n,t} \cup \mathcal{S}^+_{n,t,s}), \tag{8}$$

called a *fantasy model*. We denote the predictive mean and standard deviation function of the fantasy model $M_{n,t,s}$ as $\mu^+_{n,t,s}(\boldsymbol{x})$ and $\sigma^+_{n,t,s}(\boldsymbol{x})$, respectively. Note that $M_{n,t,s}$ is distinguished from $F^+_{n,t}$. The $\mathcal{CGP}$ $F^+_{n,t}$ can be treated as an infinite Gaussian mixture model (Rasmussen 1999), where a fantasy model $M_{n,t,s}$ is just a component of it.

With this relationship in mind, we are now ready to sample a function realization. Denote the set of fantasy models constructed by agent n in round t as $\mathcal{M}_{n,t}$. One can first sample a fantasy model $M_{n,t,s} \in \mathcal{M}_{n,t}$ and then sample a function realization $g_{n,t,s}$ from it, which gives our desired result. We have provided a visual of the process in Figure 2 in Section 4.3.

4.2.2. Acquisition Functions in Direct Expectation Form

Many famous AFs can be written in the form of expectations over function realizations. Examples include EI, NEI, PI, KG, etc. Formally, when the surrogate function is specified as $F^+_{n,t}$, there exists function $a : \mathbb{R} \to \mathbb{R}$ such that

$$\alpha(\boldsymbol{x}; F^+_{n,t}) = \mathbb{E}_{g \sim F^+_{n,t}}[a(g(\boldsymbol{x}))]. \tag{9}$$

For example, if α is specified as EI, then a is the improvement function, $a(f(\boldsymbol{x})) = \max\{\{f(\boldsymbol{x}) - y_n^{\max}\}, 0\}$ where $y_n^{\max}$ is the maximum observed response for agent n. As previously mentioned, $\alpha(\boldsymbol{x}; F^+_{n,t})$ does not have a closed-form solution as the surrogate $F^+_{n,t}$ is a $\mathcal{CGP}$. However, under a minor regularity

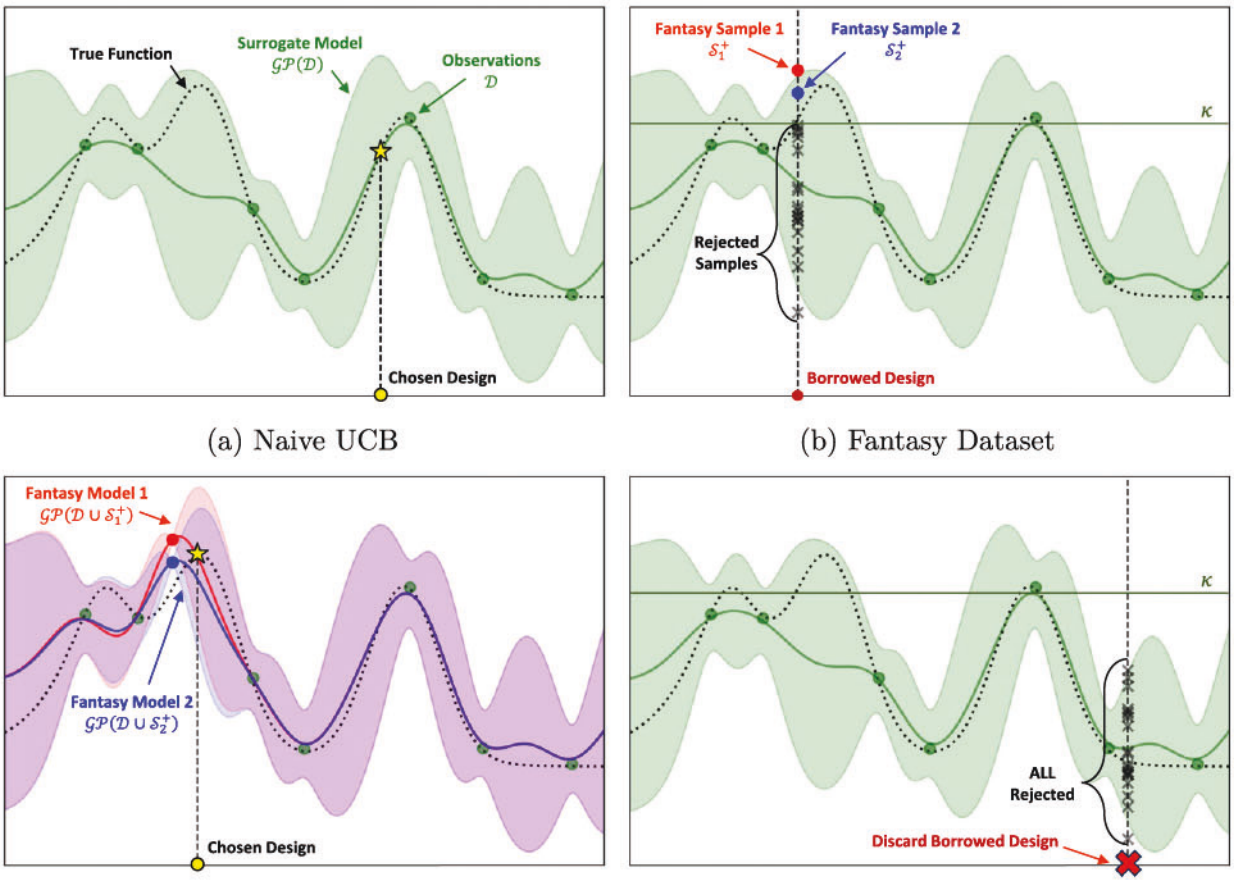

Figure 2. Illustration of $\mathcal{CGP}$-UCB. Subscripts n and t are dropped.

condition, Lemma 1 (proved in the supplementary material) gives us a way to approximate it efficiently.

Lemma 1. Assuming $\mathbb{E}_{M_{n,t,s}}\left[\alpha(x; M_{n,t,s})\right]$ exists, if there exists a such that $\alpha(x; F^+_{n,t}) = \mathbb{E}_{g\sim F^+_{n,t}}[a(g(x))]$, then

$$\alpha(x; F^+_{n,t}) = \mathbb{E}_{M_{n,t,s}}\left[\alpha(x; M_{n,t,s})\right]. \tag{10}$$

Lemma 1 simply states that $\alpha(x; F^+_{n,t})$ can be approximated by taking the expectation over the fantasy models $\mathcal{M}_{n,t,s}$ described earlier. Now, given $S_{n,t} := |\mathcal{M}_{n,t}| \le S_{\texttt{raw}}$, the AF can be approximated empirically by Monte Carlo:

$$\alpha(x; F^+_{n,t}) = \mathbb{E}_{M_{n,t,s}}\left[\alpha(x; M_{n,t,s})\right] \approx \frac{1}{S_{n,t}} \cdot \sum_{s=1}^{S_{n,t}} \alpha(x; M_{n,t,s}). \tag{11}$$

In most cases where AF is differentiable, we can get gradients $\frac{\partial}{\partial x}\alpha(x; M_{n,t,s})$. As long as $\mathbb{E}\left[\left|\frac{\partial}{\partial x}\alpha(x; M_{n,t,s})\right|\right]$ exists, we have

$$\frac{\partial}{\partial x}\alpha(x; F^+_{n,t}) = \mathbb{E}\left[\left|\frac{\partial}{\partial x}\alpha(x; M_{n,t,s})\right|\right] \approx \frac{1}{S_{n,t}} \cdot \sum_{s=1}^{S_{n,t}} \frac{\partial}{\partial x}\alpha(x; M_{n,t,s}).$$

4.2.3. Acquisition Functions with Implicit Expectation

In cases where the AF cannot be written in an expectation form, we may still be able to approximate our $\mathcal{CGP}$ AFs (Wilson et al. 2017) since many statistics are defined in their expectation form (e.g., mean, variance, entropy, etc.). Upper Confidence Bound (UCB; Srinivas et al. 2009) is a typical example in this category. It selects the surrogate UCB as the AF and its maximizer as the next design to observe. Clearly, the UCB AF requires the standard deviation. Although standard deviation itself cannot be written in the form of (9), the square of it (i.e., variance) can be approximated using Monte Carlo. Property 2 (proved in supplementary material) simplifies this procedure.

Property 2 (Fixed Covariance). For any fantasy model $M_{n,t,s}$, we can define its predictive standard deviation as $\sigma^+_{n,t} = \sigma^+_{n,t,s}$.

Property 2 simply states that the predictive variance at any design point x is the same across all the fantasy models $M_{n,t,s}$, so we can simplify the notation by dropping the s subscript. Now, denote the sample mean and variance of the predictive mean (across different fantasy models) respectively as $\bar{\mu}^+_{n,t}(\cdot)$ and $\varsigma^+_{n,t}(\cdot)$, that is,

$$\bar{\mu}^+_{n,t}(x) = \frac{1}{S_{n,t}} \cdot \sum_{s=1}^{S_{n,t}} \mu^+_{n,t,s}(x),$$
$$\varsigma^+_{n,t}(x) := \frac{1}{S_{n,t}-1} \cdot \sum_{s=1}^{S_{n,t}} (\mu^+_{n,t,s}(x) - \bar{\mu}_{n,t}(x))^2. \tag{12}$$

By the law of total variance, the $\mathcal{CGP}$-UCB can be approximated by

$$\alpha(x; F^+_{n,t}) \approx \bar{\mu}^+_{n,t}(x) + \beta_t\sqrt{\sigma^+_{n,t}(x)^2 + \varsigma^+_{n,t}(x)}, \tag{13}$$

where β_t is a user-defined hyper-parameter.

4.3. Illustration

To contextualize our approach, we start with an illustrative example that highlights the benefits of collaboration by sharing constraints. Our illustration uses UCB as the AF.

Figure 2(a) shows the result of the chosen design point at round t from agent n without collaboration. In Figure 2(b), the agent n borrows $x^+_{n',t}$ from its leader and then samples $\mathcal{S}^+_{n,t,s} \sim p(\mathcal{D}^+_{n,t}|\mathcal{D}_{n,t}, E_{n,t})$ by rejecting with respect to $\kappa_{n,t}$. Here, two fantasy datasets $\mathcal{S}^+_{n,t,1}$ and $\mathcal{S}^+_{n,t,2}$ are illustrated from which two fantasy models $M_{n,t,1} := \mathcal{GP}(\mathcal{D}_{n,t} \cup \mathcal{S}^+_{n,t,1})$ and $M_{n,t,2} := \mathcal{GP}(\mathcal{D}_{n,t} \cup \mathcal{S}^+_{n,t,2})$ are created in Figure 2(c). As shown in Figure 2(c), by incorporating the fantasy datasets $\mathcal{S}^+_{n,t,s}$, the surrogates around the borrowed point $x^+_{n',t}$ are pushed upwards. In turn, this leads to a higher upper confidence bound where the next chosen design point is by far closer to the true optimal. This illustrative example highlights the benefits of conditioning on $E_{n,t}$ to guide an agent's sampling decision toward better solutions.

Notice that in Figure 2(d), another borrowed design (call it $x^+_{n'',t}$) in $E_{n,t}$ is discarded by agent n since none of the $S_{\texttt{raw}}$ samples satisfy the constraint in (6) for both designs $x^+_{n',t}$ and $x^+_{n',t}$. Namely, $\nexists s \in [S_{\texttt{raw}}]$ such that $g_{n,t,s}(x^+_{n',t}) > \kappa_{n,t}, \forall x^+_{n',t} \in E_{n,t}$. In such cases, without taking more raw samples, agent n simply drops $x^+_{n'',t}$ from $E_{n,t}$, since none of the samples will be accepted otherwise. By doing so, agent n discards shared knowledge that significantly contradicts their posterior belief. The benefit of discarding these borrowed designs is further highlighted in Section 4.4.1. Note that although we plot only one borrowed sample for a neat presentation, $[g_{n,t,s}(x^+_{n',t})]_{x^+_{n',t}\in E_{n,t}}$ are sampled jointly from a multivariate normal distribution when there are multiple borrowed points.

4.4. Practical Guidance

4.4.1. Automatic Design Selection by Rejection Sampling

As shown in Section 4.3, some borrowed designs may be discarded by agent n if none of the samples satisfy constraint (6). This is an *intrinsic advantage* of our model that hedges against heterogeneity. In practice, due to heterogeneity across clients and various other reasons (e.g., randomness, misspecified hyperparameters, bad surrogates, etc.), we may encounter cases where some borrowed design points $x^+_{n',t} \in E_{n,t}$ violate constraints $f_n(x^+_{n',t}) > \kappa_{n,t}$, which end up having agent n conditioning on wrong information.

However, with rejection sampling, if $E_{n,t}$ conflicts agent n's surrogate $F_{n,t}$, then most of the $S_{\texttt{raw}}$ samples taken from $p(\mathcal{D}^+_{n,t}|\mathcal{D}_{n,t})$ will be rejected. We take advantage of this property and set a minimum number $S_{\texttt{quorum}}$ of required samples for the rejection sampling. When the number of retained samples is less than $S_{\texttt{quorum}}$ for a borrowed design $x^+_{n',t}$, we discard it from $E_{n,t}$. An example is shown in Figure 2(d). Therefore, all designs in $E_{n,t}$ should satisfy the minimum sample requirement. This way, we automatically delete the subset of borrowed designs that conflict with agent n's current surrogate $F_{n,t}$. As shown in our theory, due to the ability to reject borrowed misleading design points, all the fantasy datasets generated within a finite $S_{\texttt{raw}}$

enjoy concentration around the posterior mean $\mu_{n,t}(\cdot)$. This allows the algorithm to converge regardless of the heterogeneity level.

4.4.2. Scaling with Grouping Dynamics

Most of the calculations within our approach are done locally by each agent. Cross-calculations are done to establish the comparison in (2), and accordingly decide which designs to share. In a decentralized setting, pairwise comparisons between all agents can be costly. Also, more Monte Carlo samples are needed when the number of borrowed designs is large, as an accepted fantasy dataset should satisfy the constraints in (6).

To aid the scalability of our algorithm, we introduce a random grouping dynamic and conclude our algorithm in Algorithm 1. In each round, agents are evenly and randomly assigned into subgroups denoted as $\mathcal{G}_{m,t}$ such that the number of agents in the groups does not exceed $N_{\max}$. For simplicity, we denote the number of groups as $N_{\texttt{group}} := \lceil N/N_{\max} \rceil$. The algorithm is run in parallel for the groups. First, each agent fits their own $\mathcal{GP}$ using their local dataset $\mathcal{D}_{n,t}$. With the fitted $\mathcal{GP}$, agents find $x^+_{n,t+1}$ by maximizing the LCB. Next, agents communicate with each other to compare their $\kappa_{n,t}$ with the maximum LCB from other agents in the group and accordingly share the designs. Finally, each agent constructs multiple fantasy models $\mathcal{M}_{n,t,s}$ and determines the next sample point by maximizing the AF $\alpha(x; F^+_{n,t})$.

for round $t = 0, \ldots, T$ **do**
  Randomly and evenly partition agents into $N_{\texttt{group}}$ subgroups;
  for subgroup index $m = 1, \ldots, N_{\texttt{group}}$ **in parallel do**
    for agent $n \in \mathcal{G}_{m,t}$ **in parallel do**
      Fit a the Gaussian process model $\mathcal{GP}(\mathcal{D}_{n,t})$;
      Find the LCB maximizer $x^+_{n,t} \in \arg\max_{x \in \mathcal{X}} \mu_{n',t}(x) - \eta_t \cdot \sigma_{n',t}(x)$;
    end
    for agent pair $n, n' \in \mathcal{G}_{m,t}$ **in parallel do**
      if $\mu_{n',t}(x^+_{n,t}) - \eta_t \cdot \sigma_{n',t}(x^+_{n,t}) > \kappa_{n,t}$ **then**
        Agent n' send $x^+_{n',t}$ to agent n;
      end
    end
    for agent $n \in \mathcal{G}_{m,t}$ **in parallel do**
      Construct a set of fantasy models $\mathcal{M}_{n,t}$;
      Maximize the acquisition function $x_{n,t+1} \in \arg\max_x \alpha(x; F^+_{n,t})$;
      Observe the response $y_{n,t+1}$ at $x_{n,t+1}$ and augment current dataset.
    end
  end
end

Algorithm 1: Collaborative BO via Constrained $\mathcal{GP}$s

4.4.3. Cross-Communication and Synchronization

As mentioned in Section 4.1, determining whether a design $x^+_{n',t}$ qualifies to join $E_{n,t}$ involves comparing $\mu_{n',t}(x^+_{n',t}) - \eta_t \cdot \sigma_{n',t}(x^+_{n',t})$ with $\kappa_{n,t}$. This can be done in either a centralized or decentralized manner, per the need of the practical application.

Centralized. If a trusted central server is available (as assumed by previous work (Dai, Low, and Jaillet 2021)), each agent only needs to share with the central server three pieces of information, namely the LCB maximizer $x^+_{n,t}$, the value of the maximum LCB $\mu_{n',t}(x^+_{n',t}) - \eta_t \cdot \sigma_{n',t}(x^+_{n',t})$, and $\kappa_{n,t}$, which totals an upload cost of $D + 2$ scalars. Upon collecting these quantities from the agents, the server then constructs $E_{n,t}$ and sends it to agent n, which totals a download cost of at most $DN_{\max}$ scalars. This features much less communication burden compared to previous work (Dai, Low, and Jaillet 2021), which requires uploading and downloading a random Fourier feature vector.

Decentralized. Our method also works if there is no trusted central server, as the comparison can be easily done privately without raw data exchange. The secure comparison of two quantities is known as Yao's Millionaires problem (Yao 1982), which is one of the first studied and the most basic problems in secure multiparty computation (MPC, see details in section 1 of supplementary material). As such, this comparison can be done without sharing $\kappa_{n,t}$ or $\mu_{n',t}(x^+_{n',t}) - \eta_t \cdot \sigma_{n',t}(x^+_{n',t})$, and the only information agent n can receive is $E_{n,t}$.

Synchronization. Our method can be readily used in synchronous or asynchronous settings. Algorithm 1 is a synchronous algorithm where agents run a single experiment and then make comparisons in parallel to proceed to the next steps. However, our method is also applicable in offline and asynchronous settings. For instance, agent n can borrow a design from another agent n', condition on it, and proceed to find its next experiment without requiring agent n' to do the same.

We indeed envision such an asynchronous setting to be valuable when exploiting *historical databases*. For example, manufacturers may have different production schedules, but a manufacturer n can learn from others anytime since all it needs is to borrow some designs within $E_{n,t}$. A borrowed design can also be recommended by a human expert who confidently believes a certain design is good. Our methodology in such settings.

4.4.4. Space Filling Initialization

Often, in BO a good initialization can significantly improve the acquisition of better and faster solutions. To this end, we propose a sampling heuristic for initialization called Collaborative Latin Hypercube Sampling (CLHS). CLHS is a collaborative extension of the famous space-filling heuristic Latin Hypercube Sampling (LHS; McKay, Beckman, and Conover 1979). We defer the details of our initialization to section 3 in the supplementary material.

4.5. Flexibility on Selecting Synthetic Designs

As we will soon see in Section 5, due to the power of rejection sampling explained in Section 4.4.1, our theoretical analysis does *not* rely on any assumptions on the relationship between the different black box functions f_n, nor the way how an agent proposes $x^+_{n,t}$. This makes our framework highly flexible in selecting $x^+_{n,t}$, especially since understanding or defining heterogeneity for black-box functions is very challenging. In fact, regardless of how an agent selects a synthetic design $x^+_{n,t}$, our

theoretical results hold the same due to the rejection sampling scheme.

Although we gave a simple example of borrowing the LCB maximizers that are larger than one's posterior mean, one might opt for other choices due to practical considerations. In this section, we present two examples of possible alternations. Note that none of the following alternations change our theoretical analysis.

Differences on Scaling. Recall that as shown in 2, agent n' shares a design $x^+_{n',t}$ with agent n only if its LCB $\mu_{n',t}(x^+_{n',t}) - \eta_t \cdot \sigma_{n',t}(x^+_{n',t})$ is greater than $\kappa_{n,t}$. One implicit hope here (although not used in our theory) is that we expect the location where others have a high objective value will also give us a high objective value. When collaboration happens in practice, we usually expect functions f_n to have a similar scale on the response, so it is reasonable to borrow design based on the comparison above.

Yet, if one believes that the functions have different scales, then there is a simple fix whereby Agent n' shares $x^+_{n',t}$ with agent n without the comparison in 2. In other words, every agent n shares their $x^+_{n,t}$ with everyone else in their group. Still, one can rely on the rejection sampling mechanism to let agents determine which shared points they should use for conditioning. Additional experiments on this method can be found in the supplementary material.

Enhancing Privacy. Recall that one of our goals is to keep the raw data $\mathcal{D}_{n,t}$ visible to only agent n. Although we do not share any information on the response, there might be a privacy concern about sharing *synthetic* designs. We acknowledge that, because the key idea of collaboration in this article is to share informative synthetic designs, we are not aiming for Shannon's perfect secrecy by definition. That said, there are many simple methods to achieve differential privacy. As we have discussed in Section 4.1, all one needs to do is just to propose a near-optimal location guiding others toward the optimal regions. The easiest, perhaps, is adding Laplacian noise to the LCB maximizer, which guarantees differential privacy. Indeed, one can choose different sharing strategies under our framework as long as the shared point indicates a potential region of improvement. Moreover, agents can follow their own rules to propose $x^+_{n,t}$. Such a decision can be application and client-specific.

5. Theoretical Analysis

We start by acknowledging that theory in BO is still underdeveloped. Similar to previous work on federated BO (Dai, Low, and Jaillet 2020), our proof is also based on the framework of the recent theoretical advance in BO (Chowdhury and Gopalan 2017). However, a key feature of our proof is that we relax the assumption of Dai, Low, and Jaillet (2020) that heterogeneity is upper-bounded.

Now, for theoretical development, we set up some assumptions and notations.

Assumption 1. For every agent n, f_n belongs to the Reproducing kernel Hilbert space (RKHS) of real-valued functions associated with known kernel k_n where the RKHS norm is bounded by B.

Assumption 2. Noises $\boldsymbol{\varepsilon}_{n,t}$ are R-sub-Gaussian with respect to canonical filtration $\mathcal{F}_{t-1}$. Without the loss of generality, assume $\lambda \leq R \leq 1$ and $k_n(\boldsymbol{x}, \boldsymbol{x}) \leq 1$ for all $\boldsymbol{x} \in \mathcal{X}$.

This assumption is without the loss of generality because data points can be otherwise normalized to achieve the assumption.

Now to build the notation used in our theorems, we define the set of observed design points at time t as

$$X_{n,:t} = \{\boldsymbol{x}_{n,1}, \ldots, \boldsymbol{x}_{n,t}\}.$$

The corresponding responses at $X_{n,:t}$ are denoted as $\boldsymbol{y}_{n,:t} = \boldsymbol{f}_{n,:t} + \boldsymbol{\varepsilon}_{n,:t}$ where $\boldsymbol{f}_{n,:t}$ is the true underlying responses and $\boldsymbol{\varepsilon}_{n,:t}$ are the corresponding noises. We also denote the maximum information gain at time t as

$$\gamma_t := \max_{X_{n,t} \subset \mathcal{X}} \mathcal{I}(\boldsymbol{y}_{n,:t}; \boldsymbol{f}_{n,:t}),$$

where $\mathcal{I}$ denotes mutual information, and the maximum is with respect to all the possible sequences of observations $X_{n,:t}$.

The key challenge in our proof is that our fantasy model $M_{n,t,s}$ is created from both real and fantasy datasets. A major step in proving most convergence bounds in BO is to control the mean prediction error of the surrogate model. Although the deviation of real data points is limited by the sub-Gaussian Assumption 2, there is no natural restriction on the deviation of a fantasy dataset since we do not assume a known structure or upper bound on the potential heterogeneity across agents.

Fortunately, thanks to the automatic design selection scheme discussed in Section 4.4.1, we know that every sampled response in the fantasy dataset comes from the predictive distribution of $F_{n,t}$. By the generalization bound on the $F_{n,t}$, we can control the maximum deviation of these responses with respect to the ground truth using the Chernoff bound. Since we only have a finite number of borrowed designs $N_{\max}$, and the deviation of every point in the fantasy datasets is upper-bounded, the estimation error of any $M_{n,t,s}$ is limited. Therefore, we are able to obtain the bound on the generalization error of the fantasy models in Theorem 3.

Note that in order to keep generality, we do not assume how $E_{n,t}$ is generated, so this theorem holds even if borrowed design $\boldsymbol{x} \in E_{n,t}$ do not satisfy constraints $f_n(\boldsymbol{x}) > \kappa_{n,t}$.

Theorem 3 (Generalization Bound for Fantasy Models). Under Assumptions 1 and 2, letting $\hat{\lambda}_t = 1 + \frac{2}{t}$, for arbitrary $\delta \in (0, 1)$ and $T \in \mathbb{N}^+$, with a probability of at least $1 - \delta/T$, for any agent n and any sample index s, we have

$$|\mu^+_{n,t,s}(\boldsymbol{x}) - f_n(\boldsymbol{x})| \leq \sqrt{2} N_{\max} Q_{n,t} \sigma^+_{n,t}(\boldsymbol{x}),$$

where $Q_{n,t} = B + R\sqrt{2(\gamma_t + 1 + \ln(4NT/\delta))} + \sqrt{2\ln(4S_{\text{raw}}NT/\delta)}$.

Theorem 4 bounds the mean prediction error of fantasy model $M_{n,t,s}$ using its standard deviation. It implies that fantasy model $M_{n,t,s}$ is accurate with high probability wherever the predicted standard deviation is small.

With the help of Theorem 3, we can anchor the instantaneous regret on the mean prediction error, which is bounded by the predictive variance. Hence, as we have more observations, the predictive variance decays, and the regret will also decay as

desired. Now we prove the regret bound for $\mathcal{CGP}$ with upper confidence bound ($\mathcal{CGP}$-UCB) in Theorem 4. We use $\tilde{\mathcal{O}}$ to omit the logarithmic dependency on T.

One of the desired theoretical properties of BO methods is the asymptotic optimality. For agent n, denote the instantaneous regret at time t as

$$r_{n,t} = |f_n(\boldsymbol{x}_n^*) - f_n(\boldsymbol{x}_{n,t})|, \tag{14}$$

and the cumulative regret up to time T as

$$R_{n,T} = \sum_{t=1}^{T} r_{n,t}. \tag{15}$$

To show the asymptotic optimality, it suffices to show that the cumulative regret $R_{n,T}$ grows sublinearly with T, which implies that $\lim_{t\to\infty} f_n(\boldsymbol{x}_{n,t}) = f_n(\boldsymbol{x}_n^*)$.

Theorem 4 (Regret Bound for $\mathcal{CGP}$-UCB). Under all the assumptions and conditions in Theorem 3, letting $\beta_t = \sqrt{2N_{\max}Q_{n,t}}$, for any $\delta \in (0,1)$, with a probability of at least $1-\delta$, for any agent in the system, the cumulative regret $R_{n,t}$ of $\mathcal{CGP}$-UCB is upper bounded by

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

Having a cumulative regret bounded by $\tilde{\mathcal{O}}\left(\gamma_T\sqrt{T}\right)$ is standard with UCB acquisition functions. Here, γ_T is determined by the type of kernel k_n. If, for example, the kernel is a Squared Exponential Kernel, then $\gamma_T = \mathcal{O}(\log(T)^{D+1})$ (Srinivas et al. 2009). As such, the cumulative regret grows with $\tilde{\mathcal{O}}(\sqrt{T})$, so the instantaneous regret is asymptotically zero. We want to highlight that this result is without any assumptions on the heterogeneity level, meaning the regret of our algorithm will converge even if the borrowed points are provided adversarially.

6. Experiments

6.1. Comparison of Collaborative and Non-collaborative BO methods

We start by comparing the performance of eight different methods for a group of $N = 16$ collaborating agents. The compared methods are (i) $\mathcal{CGP}$-NEI (i.e., our method + NEI), (ii) $\mathcal{CGP}$-UCB (i.e., our method + UCB), (iii) $\mathcal{CGP}$-TS (i.e., our method + TS), (iv) FTSDE (Dai, Low, and Jaillet 2021), (v) FTS (Dai, Low, and Jaillet 2020), (vi) non-collaborative NEI (Letham et al. 2019), (vii) non-collaborative UCB (Srinivas et al. 2009), and (viii) non-collaborative TS (Thompson 1933). We emphasize that both FTSDE and FTS are limited to only TS, and are incompatible with other AFs like NEI. The choice of hyperparameters for FTSDE and FTS follows the guidelines in the original paper. The Matern-5/2 kernel is used for all the experiments for a fair comparison. We set $N_{\max} = 4$, $S_{\texttt{raw}} = 10^5$, and $S_{\texttt{quorum}} = 5$. Notice that $N_{\max}$ depends on the communication budget, while $S_{\texttt{raw}}$ and $S_{\texttt{quorum}}$ depend on the computation budget.

We test four famous synthetic functions of different dimensions (Balandat et al. 2020). The search spaces are set to be hypercubes denoted as $[\texttt{lb}, \texttt{ub}]^D$, where `lb` and `ub` denote lower bound and upper bound, respectively. The tested functions are (i) 5-dimensional function `Ackley-5` - $\mathcal{X} = [-32.768, 32.768]^5$, (ii) 6-dimensional function `Hartmann-6` - $\mathcal{X} = [0,1]^6$, (iii) 8-dimensional function `Cosine-8` - $\mathcal{X} = [-1,1]^8$, and (iv) 20-dimensional function `Levy-20` - $\mathcal{X} = [-10,10]^{20}$. To mimic real-world settings, we introduce heterogeneity. Specifically, we set the function f_n by shifting the standard test function f by $\boldsymbol{\psi}_n$, where $\boldsymbol{\psi}_n$ is uniformly drawn from a ball centered at 0 with a radius of $0.05|\texttt{ub} - \texttt{lb}|$, that is,

$$f_n(x) = f(\boldsymbol{x} + \boldsymbol{\psi}_n),$$
$$\text{where } \boldsymbol{\psi}_n \sim \text{Uniform}\Big(\mathcal{B}\left(0.05|\texttt{ub} - \texttt{lb}|\right)\Big).$$

Observational noise is set as $\varepsilon_{n,t} \sim \mathcal{N}(0, 0.1^2)$. We also set $\eta_t = 2$ and $\beta_t = 2$ for LCB and UCB.

We start by sampling initial designs for each client in the initialization stage to make sure the number of training data is more than the number of learnable hyperparameters in the $\mathcal{GP}$. We sample 10 initial designs for `Ackley-5`, `Hartmann-6`, and `Cosine-8`, while 25 for `Levy-20`. We use CLHS to choose the initialization for all the methods except for FTSDE since it has its own designated initialization scheme that is incompatible with CLHS. Each method on each test function is replicated 5 times. The optimality gap is defined as

$$\bar{G}_t := \frac{1}{N}\sum_{n=1}^{N}(f_n(\boldsymbol{x}_n^*) - f_n(\arg\max_{\boldsymbol{x}}\mu_n(\boldsymbol{x}))).$$

This metric measures the solution quality if the system stops after t rounds. It typically applies to offline settings like simulations where only the quality of the final solution is concerned but not the tested designs. Instantaneous regret is defined as

$$\bar{r}_t := \frac{1}{N}\sum_{n=1}^{N}(r_{n,t}).$$

This metric measures the solution quality at each round t. It typically applies to settings like business operations where the reward of all the tested designs is important.

The results, in logarithmic scales, are shown in Figure 3. The solid lines represent the means of the five replications. The upper bound of the shaded region denotes the maximum value of the five replications, and the lower bound denotes the minimum value. From the figures, we find interesting insights. First, collaborative methods provide a clear improvement over their non-collaborative counterparts. Second, our $\mathcal{CGP}$ approaches show consistent superior performance across the benchmarks. Third, we notice the importance of adaptability to different AFs. For example, even non-collaborative UCB and NEI outperform FTSDE and LTS on `Ackley5` and `Cosine8` functions, resulting in an even more significant competitive edge of $\mathcal{CGP}$-UCB and $\mathcal{CGP}$-NEI over FTSDE. Additionally, the results showcase significant variations in both the optimality gap and instantaneous regret. AFs with a higher emphasis on exploration tend to display a smaller optimality gap, whereas those prioritizing exploitation tend to exhibit smaller instantaneous regret. For

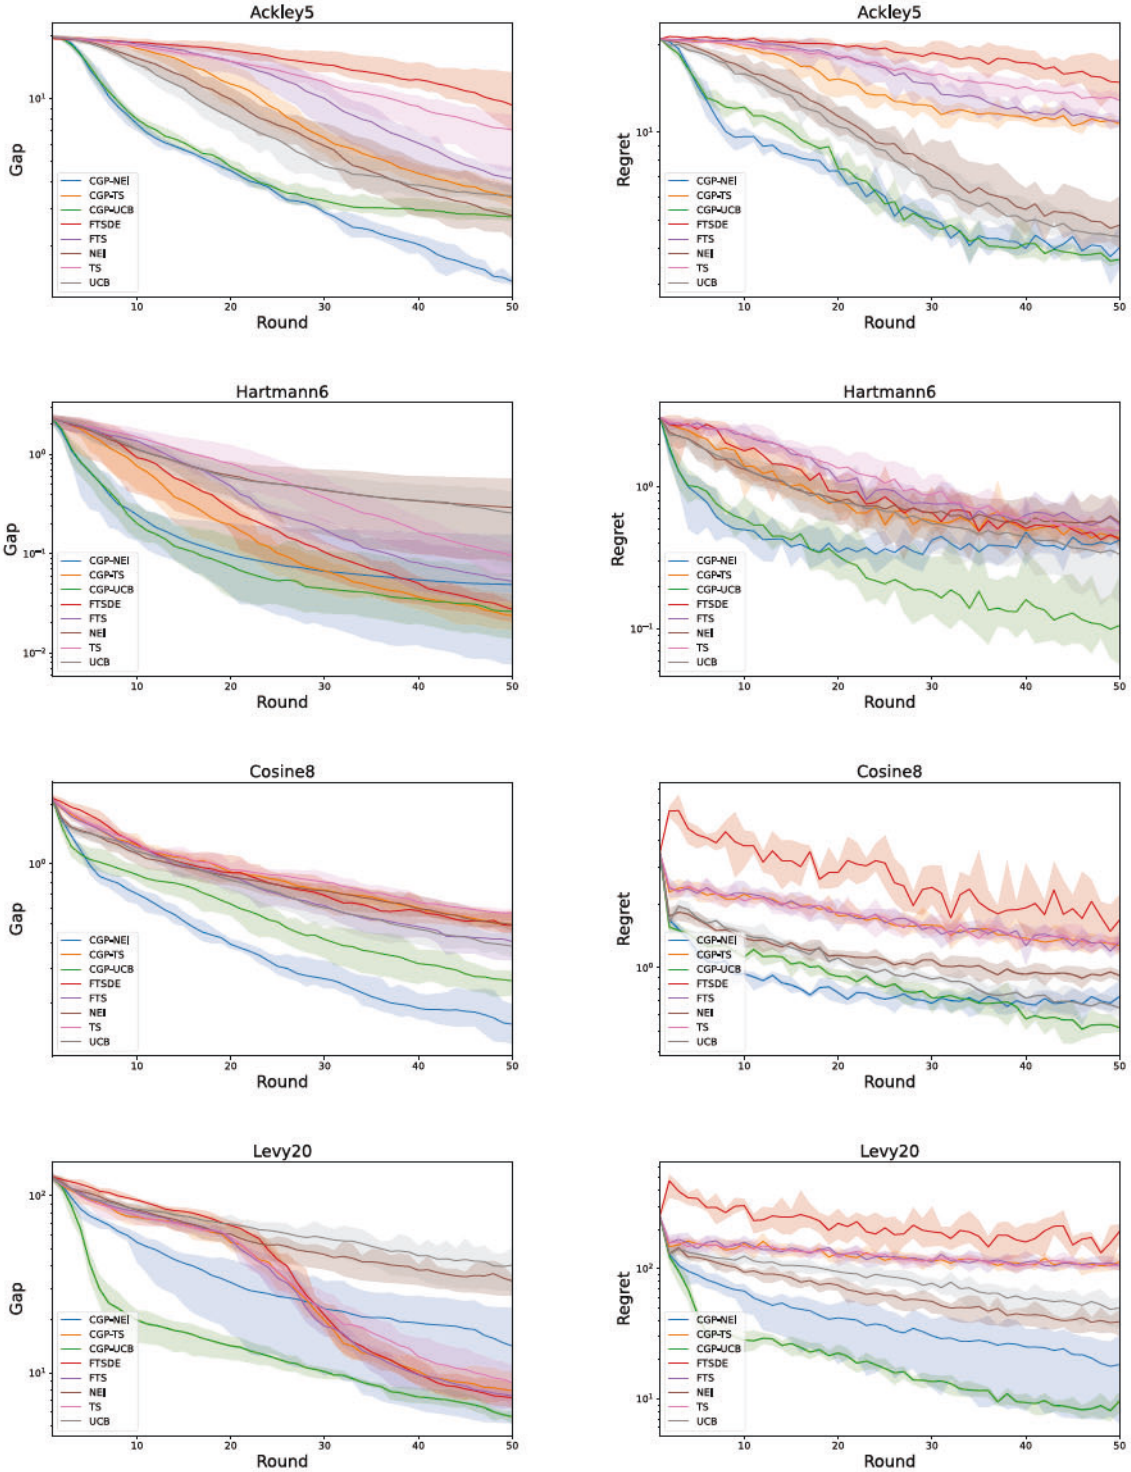

Figure 3. Optimality gap (left) and instantaneous regret (right) of four test functions.

example, one can observe that UCB features more exploitation than exploration compared to TS, which gives it a smaller instantaneous regret and makes it a better option for online implementations. As mentioned in Section 2, the selection of AFs is application-specific. Our experiment echoes this statement and highlights the importance of collaborative methods to be compatible with a large variety of AFs, which is a key contribution of this work.

6.2. The Power of Collaboration

We also test the benefits gained from more collaborators. We use the `Levy-20` function with a UCB AF for $N \in \{1, 4, 16, 64, 256\}$ collaborating agents. Other settings are kept the same as before. The results from our $\mathcal{CGP}$-UCB approach are shown in Figure 4. From the figure, one can directly observe that both optimality gaps and instantaneous regrets decay faster

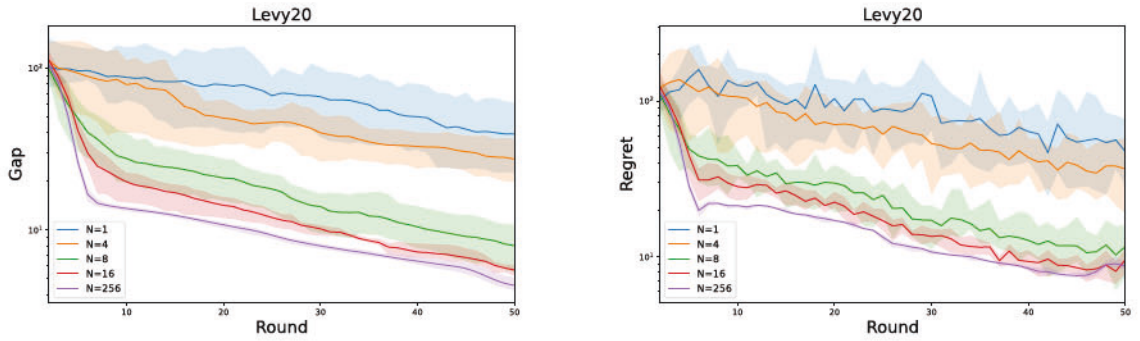

Figure 4. The power of collaboration.

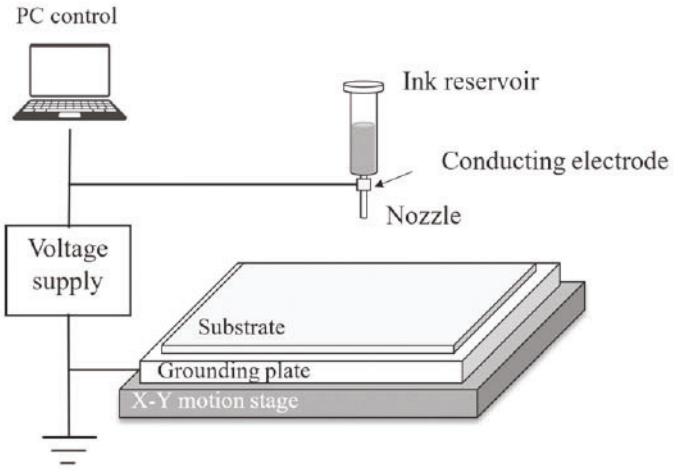

Figure 5. EHD inkjet printing setup.

as the number of collaborators increases. Yet this power of collaboration has a diminishing marginal rate of return.

7. Electrohydrodynamic Inkjet Printing

In addition to the numerical studies above, we conducted a limited-budget case study. Our goal is to showcase the benefits of collaboration for the process parameter tuning of a 3D printing process. The process used, known as Electrohydrodynamic (EHD) inkjet printing, is a patterning method using electric fields to generate fluid flows that deliver materials to a substrate. The printing system (shown in Figure 5) consists of an ink reservoir, a printhead, a substrate, a grounding plate, a power supply, and a three-axis (XYZ) stage. The ink reservoir is a container that holds the ink to be printed. The printhead contains a nozzle and a conducting electrode, which connects the power supply and the metal nozzle. The grounding plate is a conducting plate used as the counter electrode, known as the ground electrode. In the printing process, a high-voltage signal is applied between the nozzle and the ground electrode, creating a strong electric field, and forming a conical shape at the tip of the nozzle. When the electric field is strong enough, the conical shape will break up, and droplets will be emitted from the tip of the cone shape. The nozzle is attached to the Z axis, which is able to move up and down, and the X–Y stages provide movement in the X–Y directions during the printing. In this experiment, a 75mm × 25mm × 1mm glass substrate, a 100μm glass nozzle, and a pulse voltage signal were used.

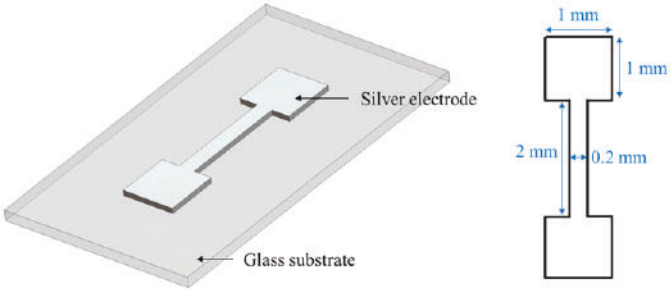

Figure 6. Silver electrode.

In this case study, a silver electrode was fabricated, which consists of two 1mm-by-1mm solid pads at two sides and a bridge with a 2 mm length and 0.2 mm width in the center, as shown in Figure 6. Our goal was to find the best set of process parameters that minimize the production time for fabricating a single silver electrode.

The production time of a designed pattern is influenced by factors such as the flow rate of ink deposition, the pattern design, and the printing speed. These, in turn, are affected by a combination of parameters, including voltage potential, duty ratio, pulse frequency, nozzle size, stand-off distance, plotting speed, and ink properties. In a typical optimization process, the system and the ink are usually not changed, so the geometry parameters (nozzle size, stand-off distance) and the ink properties are usually determined ahead of the parameter optimization process. Hence, the four process parameters to be optimized in this case study are voltage potential, pulse frequency, duty ratio, and plotting speed.

7.1. Experimental Setting

The production time is defined as the time to complete a single electrode printing in this case study. The feasible range of process parameters ($\mathcal{X}$) was established based on previous printing tests on the same printing system.

Voltage is one of the most effective factors in controlling the droplet size. However, it must exceed the onset voltage to initiate printing and stay below a certain threshold to prevent unstable jetting status, such as multi-jet status. The duty ratio and the frequency of the pulse signal contribute to both droplet size and jetting frequency. They need to be carefully selected to match the

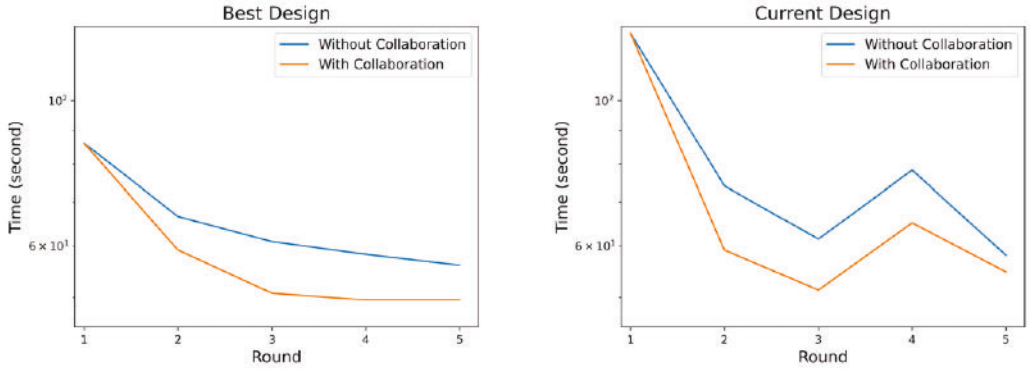

Figure 7. Optimality gap (left) and instantaneous regret (right) versus round.

plotting speed to achieve solid, continuous traces. If the plotting speed is too low or high, the printed trace will become separate dots or bulge lines. Thus, the feasible region $\mathcal{X}$ was determined with voltage ranging from 2000 to 3000V, duty ratio from 20 to 30, frequency from 10 to 100Hz, and plotting speed from 0.4 mm/s to 2 mm/s.

Our experiment featured both collaborative and non-collaborative groups, each consisting of $N = 3$ trained agents conducting the experiment. Due to budget constraints, the agents ran the experiments over $T = 5$ rounds, with the first round serving as an initialization round. In this round, each agent selected five designs using CLHS. To ensure a fair comparison, both the collaborative and non-collaborative groups shared the same set of initial designs. In each experiment, agents tested the quality of the printing parameters by conducting experiments with their chosen printing settings and recording the production time. Each experiment took approximately 70 min to complete, including the setup time and nozzle preparation. We used the NEI AF for both groups and set $\eta_t = 0$ and $N_{\max} = 3$ (since we only had three participants). Note that our experiment's objective was to showcase the benefits of collaboration. We did not impose heterogeneity across the agents due to budget constraints.

7.2. Experimental Results

The comparison between collaborative and non-collaborative methods is depicted in Figure 7. Given the real-world nature of this experiment, the ground truth is not available. Therefore, we present the following two statistics: (i) Average production time of the *best observed design* at round t across the three agents: $\frac{1}{3}y_{1,t}^{\max} + \frac{1}{3}y_{2,t}^{\max} + \frac{1}{3}y_{3,t}^{\max}$, and (ii) Average production time of the *current design* observed at round t: $\frac{1}{3}y_{1,t} + \frac{1}{3}y_{2,t} + \frac{1}{3}y_{3,t}$. These two metrics are designed to respectively resemble the optimality gap and instantaneous regret. We also provide a visualization in Figure 8 to highlight the chosen design points across the four process parameters and their respective production time. The numbers in the dots denote the corresponding round they are selected. Note that for better visualization, we only plot the best design (with the shortest time) across the 5 initialized points for each agent.

Based on Figure 7, we observe that our collaborative approach yielded significantly superior results compared to the non-collaborative counterpart in terms of both the best design and current design. Indeed, after five rounds, the collaborative group had an average best design that resulted in around 15% reduced production time compared to the best design from the non-collaborative group.

Furthermore, based on Figure 8, we observe that the collaborative group helped each other narrow down the exploration space to regions with smaller production times, where most of the testing happened (see the concentration of blue dots), whereas the non-collaborative group spent a lot of time in areas that yielded much larger production times.

8. Conclusion

In this study, we propose the first general-purpose framework for Collaborative BO using $\mathcal{CGP}$, a method that significantly expedites the trial-and-error process for black-box optimization under collaboration. In comparison to existing methods, our approach offers flexibility by allowing the use of a large class of acquisition functions and can effectively handle heterogeneity across the participating agents.

Given the relative infancy of the collaborative trial and error field, we envision that many future questions pose promising directions for future work. On the privacy side, one may enhance privacy by perturbing the shared designs $x_{n,t}^{+}$ with some disturbance, such as Laplacian noise. The tradeoff is that potentially fewer designs will be accepted by other agents, and those accepted will be less informative. Additionally, it may be worthwhile to sequentially understand the nature of the heterogeneity across agents and design techniques that adapt the shared designs accordingly.

Supplementary Materials

The supplementary materials contain (i) an explanation of evaluating comparison functions with Secure Multiparty Computation (MPC), (ii) proof of Theorems, (iii) a description of Collaborative Latin Hypercube Design (CLHS), (iv) an extended discussion on scaling, and (vi) the code for replicating experiments in this article. The experiments are implemented in Python using BoTorch (Balandat et al. 2020) with CUDA support.

Disclosure Statement

The authors report there are no competing interests to declare.

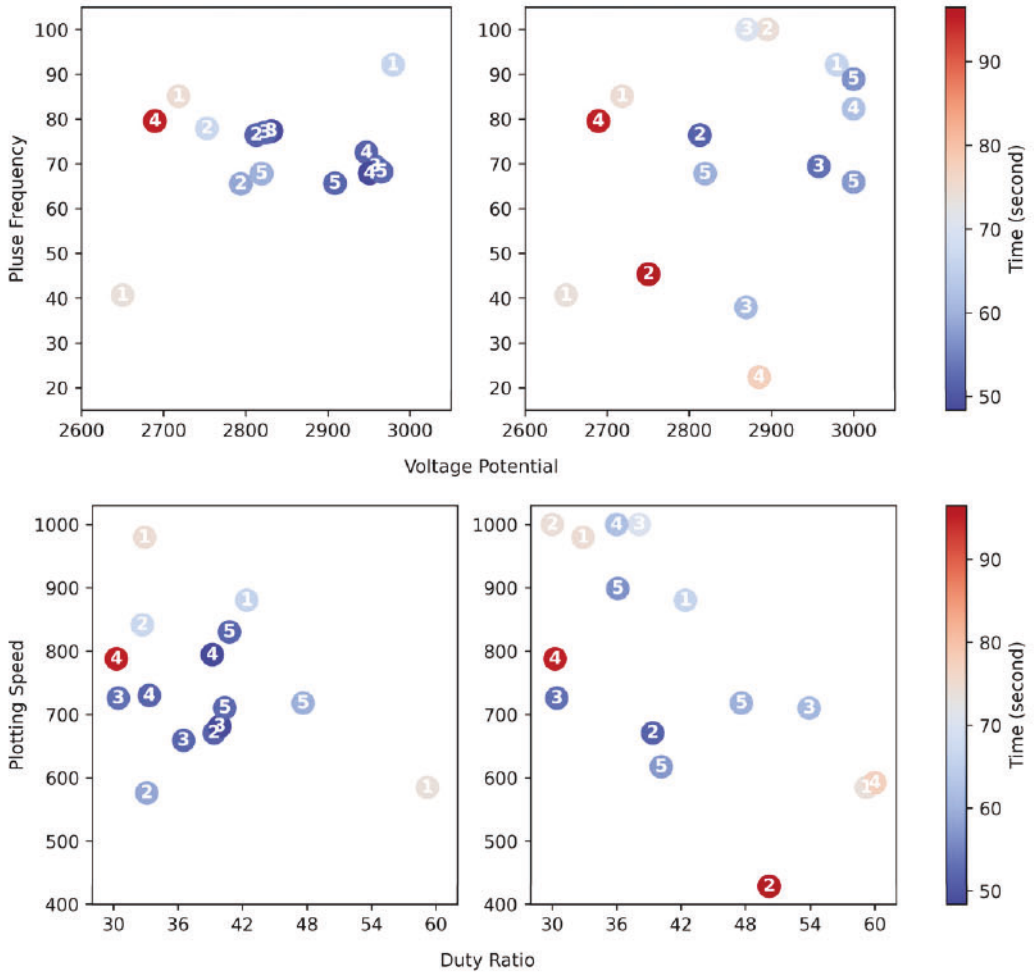

Figure 8. Designs selected by collaborative (left) & non-collaborative (right) method.

ORCID

Hantang Qin http://orcid.org/0000-0003-4180-7911
Raed Al Kontar http://orcid.org/0000-0002-4546-324X

References

Attia, P. M., Grover, A., Jin, N., Severson, K. A., Markov, T. M., Liao, Y.-H., Chen, M. H., Cheong, B., Perkins, N., Yang, Z., et al. (2020), "Closed-Loop Optimization of Fast-Charging Protocols for Batteries with Machine Learning," *Nature*, 578, 397–402. [32]

Azimi, J., Fern, A., and Fern, X. (2010), "Batch Bayesian Optimization via Simulation Matching," in *Advances in Neural Information Processing Systems* (Vol. 23). [33]

Bai, T., Li, Y., Shen, Y., Zhang, X., Zhang, W., and Cui, B. (2023), "Transfer Learning for Bayesian Optimization: A Survey," arXiv preprint arXiv:2302.05927. [33]

Balandat, M., Karrer, B., Jiang, D., Daulton, S., Letham, B., Wilson, A. G., and Bakshy, E. (2020), "Botorch: A Framework for Efficient Monte-Carlo Bayesian Optimization," in *Advances in Neural Information Processing Systems* (Vol. 33), pp. 21524–21538. [40,43]

Chowdhury, S. R., and Gopalan, A. (2017), "On Kernelized Multi-Armed Bandits," in *International Conference on Machine Learning*, pp. 844–853, PMLR. [39]

Dai, Z., Low, B. K. H., and Jaillet, P. (2020), "Federated Byesian Optimization via Thompson Sampling," in *Advances in Neural Information Processing Systems* (Vol. 33), pp. 9687–9699. [33,39,40]

——— (2021), "Differentially Private Federated Bayesian Optimization with Distributed Exploration," in *Advances in Neural Information Processing Systems* (Vol. 34), pp. 9125–9139. [33,38,40]

Daulton, S., Cakmak, S., Balandat, M., Osborne, M. A., Zhou, E., and Bakshy, E. (2022), "Robust Multi-Objective Bayesian Optimization Under Input Noise," in *International Conference on Machine Learning*, pp. 4831–4866, PMLR. [33]

Duan, W., Ankenman, B. E., Sanchez, S. M., and Sanchez, P. J. (2017), "Sliced Full Factorial-based Latin Hypercube Designs as a Framework for a Batch Sequential Design Algorithm," *Technometrics*, 59, 11–22. [33]

Duris, J., Kennedy, D., Hanuka, A., Shtalenkova, J., Edelen, A., Baxevanis, P., Egger, A., Cope, T., McIntire, M., Ermon, S., et al. (2020), "Bayesian Optimization of a Free-Electron Laser," *Physical Review Letters*, 124, 124801. [32]

Friedman, M., and Savage, L. (1947), "Planning Experiments Seeking Maxima," in *Selected Techniques of Statistical Analysis for Scientific and Industrial Research, and Production and Management Engineering*, pp. 363–372. [32,33,34]

Gongora, A. E., Xu, B., Perry, W., Okoye, C., Riley, P., Reyes, K. G., Morgan, E. F., and Brown, K. A. (2020), "A Bayesian Experimental Autonomous Researcher for Mechanical Design," *Science Advances*, 6, eaaz1708. [32]

Gonzalez, J., Longworth, J., James, D. C., and Lawrence, N. D. (2015), "Bayesian Optimization for Synthetic Gene Design," arXiv preprint arXiv:1505.01627. [32]

Gramacy, R. B. (2020), *Surrogates: Gaussian Process Modeling, Design, and Optimization for the Applied Sciences*, Boca Raton, FL: CRC Press. [33]

Hotelling, H. (1941), "Experimental Determination of the Maximum of a Function," *The Annals of Mathematical Statistics*, 12, 20–45. [33]

Hunt, N. (2020), "Batch Bayesian Optimization," PhD thesis, Massachusetts Institute of Technology. [33]

Irshad, F., Karsch, S., and Döpp, A. (2021), "Expected Hypervolume Improvement for Simultaneous Multi-Objective and Multi-Fidelity Optimization," arXiv preprint arXiv:2112.13901. [33]

Jones, D. R., Schonlau, M., and Welch, W. J. (1998), "Efficient Global Optimization of Expensive Black-Box Functions," *Journal of Global Optimization*, 13, 455–492. [33]

Kontar, R., Shi, N., Yue, X., Chung, S., Byon, E., Chowdhury, M., Jin, J., Kontar, W., Masoud, N., Nouiehed, M., et al. (2021), "The Internet of Federated Things (ioft)," *IEEE Access*, 9, 156071–156113. [32,33]

Kushner, H. J. (1963), "A New Method of Locating the Maximum Point of an Arbitrary Multipeak Curve in the Presence of Noise," in *Joint Automatic Control Conference*, pp. 69–79. [33,34]

Letham, B., Karrer, B., Ottoni, G., and Bakshy, E. (2019), "Constrained Bayesian Optimization with Noisy Experiments," *Bayesian Analysis*, 14, 495–519 [33,40]

Matheron, G. (1963), "Principles of Geostatistics," *Economic Geology*, 58, 1246–1266. [34]

McKay, M. D., Beckman, R. J., and Conover, W. J. (1979), "A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code," *Technometrics*, 21, 239–245. [38]

Moriconi, R., Deisenroth, M. P., and Sesh Kumar, K. (2020), "High-Dimensional Bayesian Optimization Using Low-Dimensional Feature Spaces," *Machine Learning*, 109, 1925–1943. [33]

Raiffa, H., and Schlaifer, R. (1961), *Applied Statistical Decision Theory*, Boston: Harvard University. [33]

Rasmussen, C. (1999), "The Infinite Gaussian Mixture Model," in *Advances in Neural Information Processing Systems* (Vol. 12). [36]

Shields, B. J., Stevens, J., Li, J., Parasram, M., Damani, F., Alvarado, J. I. M., Janey, J. M., Adams, R. P., and Doyle, A. G. (2021), "Bayesian Reaction Optimization as a Tool for Chemical Synthesis," *Nature*, 590, 89–96. [32]

Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M. (2009), "Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design," arXiv preprint arXiv:0912.3995. [33,37,40]

Thompson, W. R. (1933), "On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples," *Biometrika*, 25, 285–294. [33,35,40]

Wang, Z., and Jegelka, S. (2017), "Max-Value Entropy Search for Efficient Bayesian Optimization," in *International Conference on Machine Learning*, pp. 3627–3635, PMLR. [33]

Wilson, J. T., Moriconi, R., Hutter, F., and Deisenroth, M. P. (2017), "The Reparameterization Trick for Acquisition Functions," arXiv preprint arXiv:1712.00424. [37]

Wu, J., and Frazier, P. (2016), "The Parallel Knowledge Gradient Method for Batch Bayesian Optimization," in *Advances in Neural Information Processing Systems* (Vol. 29). [33]

Yao, A. C. (1982), "Protocols for Secure Computations," in *23rd Annual Symposium on Foundations of Computer Science (SFCS 1982)*, pp. 160–164, IEEE. [38]