

Collaborative and Federated Black-box Optimization: A Bayesian Optimization Perspective

Raed Al Kontar
Industrial & Operations Engineering
University of Michigan
alkontar@umich.edu

Abstract—We focus on collaborative and federated black-box optimization (BBOpt), where agents optimize their heterogeneous black-box functions through collaborative sequential experimentation. From a Bayesian optimization perspective, we address the fundamental challenges of distributed experimentation, heterogeneity, and privacy within BBOpt, and propose three unifying frameworks to tackle these issues: (i) a global framework where experiments are centrally coordinated, (ii) a local framework that allows agents to make decisions conditioned on shared information, and (iii) a predictive framework that enhances local surrogates through collaboration to improve decision-making. We categorize existing methods within these frameworks and highlight key open questions to unlock the full potential of federated BBOpt. Our overarching goal is to shift federated learning from its predominantly descriptive/predictive paradigm to a prescriptive one, particularly in the context of BBOpt—an inherently sequential decision-making problem.

Index Terms—Collaboration, Federated, Personalization, Privacy, Heterogeneity, Experimentation, Bayesian Optimization

I. INTRODUCTION

The tremendous increase in computational capabilities of edge devices, along with the rapid market infiltration of powerful AI chips, has led to explosive interest in collaborative and distributed analytics, such as federated learning (FL), which distributes model learning across diverse and often heterogeneous data sources to process more of the user's data at its point of origin. FL addresses many of the privacy concerns, regulatory constraints, communication costs, and skyrocketing data volumes that have made traditional cloud-centric computation increasingly unsustainable.

Significant progress has been made in FL. Methods have been proposed to enable faster convergence [1, 2], address heterogeneity in size and distribution [3, 4], improve parameter aggregation schemes [5], personalize across concept and covariate shifts [6, 7], protect against adversarial attacks [8], promote fairness [9, 10], and quantify uncertainty [11, 12], among many others (see [13] for a detailed review). To date, these efforts have focused mainly on *predictive* modeling, where the goal is to create a global or personalized predictive map (often a deep network) that leverages knowledge from different sources while circumventing the need to share raw data. In addition, recent *descriptive* FL literature has been proposed to better understand the shared and unique features

across diverse datasets. This work focuses on distributed low-rank decomposition methods such as PCA [14], matrix completion [15, 16], and dictionary learning [17].

Yet, a key opportunity lies in advancing FL from a primarily predictive/descriptive paradigm to a *prescriptive* one, which remains in its infancy. While many avenues exist for exploration, this paper provides a vision and mathematical framework for prescriptive FL in the context of black-box optimization (BBOpt) —*an inherently sequential decision making problem*. Needless to say, the success of many real-world problems critically depends on trial & error, where the goal is to manipulate a set of variables, hereon referred to as *designs*, to achieve an optimal outcome. At its core, BBOpt represents the mathematical framework for trial & error, where the relationship between designs and outcomes is often unknown (i.e., black-box), and optimal designs can only be identified through sequential experimentation. Since experiments are often expensive, and search regions can be high-dimensional, the goal of BBOpt is to carefully decide on the next-to-observe design(s), in order to find a good design with the fewest trials possible.

Now, if a fleet of agents exists, federated BBOpt sets forth a collaborative approach whereby agents sequentially distribute their experimentation efforts to improve and fast-track their optimal design process. If successful, federated BBOpt can significantly reduce trial & error cost and time and benefit all participating entities, all while circumventing the need to share raw data to preserve privacy, security, and intellectual property. Despite its appeal, fundamental challenges must be first addressed to enable federated BBOpt.

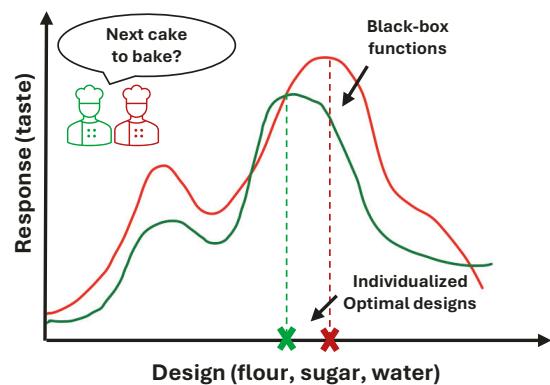


Fig. 1. Collaborative & federated black-box optimization

This work is supported by NSF CAREER Award 2144147 and the NSF CMMI-FMRG Award 2328010

- **Challenge 1:** How to effectively distribute experimentation across collaborating entities?
- **Challenge 2:** How to collaborate in the presence of heterogeneity? Heterogeneity can be both in the black-box function across agents, as well as in their resources and fidelity. Here, it is worthwhile noting that if agents are homogeneous, solutions do exist from the rich literature on Batch BBOpt [18, 19].
- **Challenge 3:** How to design collaboration that respects the privacy of all collaborating entities? This may be critical to persuade agents to join the collaborative process.

Fig. 1 provides an anecdotal example, where two chefs collaborate to find the optimal flour, sugar and water levels for baking the best cake. While the design-response relationships share commonalities, the chefs may have slightly different palates, and they may only collaborate if their recipes and outcomes remain private. Given these challenges and the emerging nature of this field, this paper sheds light on mathematical frameworks that enable federated BBOpt and highlights key open questions that must be addressed to fully unlock its potential.

II. SETTING THE STAGE

In the literature, BBOpt has been approached from multiple angles, including Bayesian optimization (BO), derivative-free optimization, and evolutionary algorithms. While all these methods have the potential to be extended to federated settings, we primarily focus on BO [20]. Notably, the few existing works on federated BBOpt (which will be highlighted as we proceed) mainly fall within the realm of BO.

Mathematically, the goal of BO is to find a design that optimizes a black-box function $f : \mathbb{R}^d \rightarrow \mathbb{R}$,

$$\mathbf{x}^* = \underset{\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^d}{\operatorname{argmax}} f(\mathbf{x}), \quad (1)$$

that models the underlying true relationship between a design point $\mathbf{x} \in \mathbb{R}^d$ and the response f . Clearly, since f is unknown, using first or second order optimization algorithms is not feasible, as we can only observe a potentially noisy version $y(\mathbf{x}) \triangleq f(\mathbf{x}) + \epsilon$ of f by running an experiment at \mathbf{x} . These experiments, whether through physical experimentation or simulations, come with time and budget costs. As such, one needs to carefully decide on the next-to-observe design. To do so, BO resorts to a utility function $U(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}$ [21] that quantifies the benefits gained if one were to conduct an experiment at design point \mathbf{x} .

A common example is the improvement utility function $U(\mathbf{x}) = \max(f(\mathbf{x}) - y^*, 0)$ [22] where y^* is the current best observed response. Basically, $U(\mathbf{x})$ only gives utility to design points that give a better outcome than the current best y^* . Clearly, before doing an experiment, one cannot calculate the utility as we do not know $f(\mathbf{x})$. Yet, at some time t one can start with a small initial dataset $\mathcal{D}_t = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_t, y_t)\}$ with t observations. Using \mathcal{D}_t , a surrogate model $\mathbb{P}_{f|\mathcal{D}_t}$ that estimates the relationship between f and \mathbf{x} can be built to

predict the outcome at unobserved designs. With this, the next-to-observe design at time t can be chosen as the design that gives the best utility in expectation:

$$\mathbf{x}_t^{\text{new}} = \underset{\mathbf{x}}{\operatorname{argmax}} F_t \triangleq \mathbb{E}_{\mathbb{P}_{f|\mathcal{D}_t}} [U(f(\mathbf{x}); \mathcal{D}_t)]. \quad (2)$$

Notice that expectation is taken in (2) as the surrogate $\mathbb{P}_{f|\mathcal{D}_t}$ is usually a posterior belief. Predominantly, BO surrogates are Gaussian processes (\mathcal{GP}) or more recently Bayesian deep neural networks.

III. FEDERATED FRAMEWORKS

Now, given K collaborating agents, each with their own dataset $\mathcal{D}_{k,t}$ at some time t , we assume (for now) that each agent has a budget of T experiments across T iterations. Additionally, agents can communicate with each other directly or via a central orchestrator. In this collaborative setting, our goal is to enable agents to borrow strength from one another in deciding on their individualized next-to-observe designs, $\mathbf{x}_{k,t}^{\text{new}}$.

Next, we present three unifying frameworks that enable federated BO, categorize the existing methods within these frameworks, and discuss open questions associated with each. We note that the methods introduced aim to identify the optimal $\mathbf{x}_{k,t}^{\text{new}}$ while preserving privacy and are repeated for all T iterations until the experimentation budget is exhausted.

A. Global Decisions

The first framework is one where the next experiments $\mathbf{x}_{k,t}^{\text{new}}$ are dictated by a central entity that aggregates summary statistics from all agents. Perhaps a natural start here is to follow the building-block literature in predictive FL, where the goal is to minimize the expected loss across all agents. In a federated BO setting, this translates to maximizing an average over all the agents expected utility functions $F_{k,t} = \mathbb{E}_{\mathbb{P}_{f_k|\mathcal{D}_{k,t}}} [U(f_k(\mathbf{x}); \mathcal{D}_{k,t})]$ at each time t . That is:

$$\max_{\mathbf{x}} \sum_{k=1}^K p_k F_{k,t} = \sum_{k=1}^K p_k \left[\mathbb{E}_{\mathbb{P}_{f_k|\mathcal{D}_{k,t}}} [U(f_k(\mathbf{x}); \mathcal{D}_{k,t})] \right], \quad (3)$$

where p_k is some weight given to an agent $k \in [K]$, with $\sum p_k = 1$. If this was feasible, one could use the rich literature on global predictive FL, starting from the famous FedAvg algorithm [23], where at each communication round $r \in [R]$, agents perform few optimization iterates on $F_{k,t}$ to obtain local candidate solutions $\mathbf{x}_{k,t}^{c,(r)} \leftarrow \text{agent_update}(F_{k,t})$, which are then averaged centrally to get a *global* design $\mathbf{x}_t^{\text{new},(r)} = \sum_k p_k \mathbf{x}_{k,t}^{c,(r)}$. The global design then serves as an initialization to the next optimization round, until reaching $\mathbf{x}_t^{\text{new}} = \mathbf{x}_t^{\text{new},(R)}$. Then all agents run the experiment at $\mathbf{x}_t^{\text{new}}$.

However, in contrast to predictive FL, where the global design serves as an initialization for the next round, in our prescriptive setting, $\mathbf{x}_t^{\text{new}}$ represents the next-to-observe design. Consequently, having all agents run the same experiment is inefficient and leads to a suboptimal use of collaborative resources. Instead, the collaborative process should allow agents to decide on their individualized tests $\mathbf{x}_{k,t}^{\text{new}}$.

1) *Consensus*: To this end, one way to rethink federated BO is through consensus, proposed in [24], where each agent's next experiment $\mathbf{x}_{k,t}^{\text{new}}$ is a weighted combination of all agents' (including their own) candidate solutions $\{\mathbf{x}_{k,t}^c\}_{k \in [K]}$. Mathematically, this is

$$\mathbf{x}_{k,t}^{\text{new}} = [(\mathbf{W}^{(t)} \otimes \mathbf{I}_d) \mathbf{x}_t^c]_k \quad (4)$$

$$\text{such that } \mathbf{x}_{k,t}^c = [\mathbf{x}_t^c]_k = \underset{\mathbf{x}_k}{\text{argmax}} F_{k,t}, \quad (5)$$

where $\mathbf{x}_t^c = [\mathbf{x}_{1,t}^{c\top}, \dots, \mathbf{x}_{K,t}^{c\top}]^\top$ is the concatenation of “candidate” solutions that maximize each agent's individual expected utility. In addition, the consensus matrix $\mathbf{W}^{(t)}$ is a symmetric, doubly stochastic matrix (i.e., $\sum_k w_{kj}^{(t)} = \sum_j w_{kj}^{(t)} = 1$ for $j, k \in [K]$) with non-negative elements $w_{kj}^{(t)} \leq 1$, \mathbf{I}_d is a $d \times d$ identity matrix, and \otimes denotes the Kronecker product operation, i.e., $\mathbf{W} \otimes \mathbf{I}_d$ results in a $dK \times dK$ matrix.

The main idea of the formulation is intuitive. Agents first find their own utility maximizers $\mathbf{x}_{k,t}^c$, but their actual actions (experiments) $\mathbf{x}_{k,t}^{\text{new}}$ are a weighted combination of $\mathbf{x}_{k,t}^c$ from all agents in the system. The consensus matrix $\mathbf{W}^{(t)}$ dictates such weights and accordingly dictates how much one agent's decision will depend on others. Indeed, such an approach has several interesting features. (i) First, the consensus step $(\mathbf{W}^{(t)} \otimes \mathbf{I}_d) \mathbf{x}_t^c$ naturally yields K *entity-specific* designs for agents to *collaboratively explore and exploit* the search space. (ii) Second, the consensus step only requires each agent sharing $\mathbf{x}_{k,t}^c$; hence, information sharing is reduced to sharing designs without disclosing any agent-specific outcomes. (iii) Third, and perhaps most importantly, the consensus matrix $\mathbf{W}^{(t)}$ can be time-varying, allowing one to dynamically adjust the dependence of one agent's decision on another. This adds a critical layer of flexibility and enables to account for heterogeneity, which will be discussed shortly. (iv) Fourth, unlike in predictive FL (as in FedAvg), there are no iterations over R communication rounds to find $\mathbf{x}_{k,t}^{\text{new}}$. Instead, agents try to find their local optimal solutions $\mathbf{x}_{k,t}^c$ once and then send them to the cloud.

These properties above yield a naturally distributed algorithm.

Algorithm 1 Consensus iterated over T Experiments

- 1: **Trial**: Agents observe $\mathbf{y}_{k,t-1}^{\text{new}}$ at $\mathbf{x}_{k,t-1}^{\text{new}}$
- 2: **Posterior update**: Agents find $\mathbb{P}_{\hat{f}|\mathcal{D}_t}$ with new data
- 3: **Optimize**: Agents find $\mathbf{x}_{k,t}^c$ using (5) and send to cloud
- 4: **Consensus**: Cloud finds $\mathbf{x}_{k,t}^{\text{new}}$ using (4) and sends to agents

A central property of consensus is the dynamic flexibility of $\mathbf{W}^{(t)}$. Such a property is *not only helpful but necessary* when heterogeneity exists. Specifically, when $f_1 \neq \dots \neq f_K$, it is required for $\mathbf{W}^{(t)} \rightarrow \mathbf{I}$. To see this, consider Fig. 1 and assume the two agents have reached their optimal design. Then, any $\mathbf{W} \neq \mathbf{I}$ will always cause them to move away from their optimum. This inspires a simple framework in (6) where the off-diagonal elements of $\mathbf{W}^{(t)}$ decay linearly to zero and diagonal elements to 1. A similar concept for

dithering collaboration is also used in [25, 26], albeit within a different framework (introduced shortly). The intuition is as follows: In early stages, agent k lacks enough data to build a high-quality surrogate and therefore should leverage information from others. As agent k gathers more data in the later stages, they will focus more on their own objectives to find personalized optimal designs.

$$\mathbf{W}^{(t+1)} = \mathbf{W}^{(t)} + \begin{bmatrix} \frac{K-1}{TK} & -\frac{1}{TK} & \dots & -\frac{1}{TK} \\ \vdots & \vdots & \vdots & \vdots \\ -\frac{1}{TK} & -\frac{1}{TK} & \dots & \frac{K-1}{TK} \end{bmatrix}, \quad (6)$$

It should be noted that this is just one example; *many other interesting ideas can be incorporated*. For instance, one could design $\mathbf{W}^{(t)}$ for agents to be primarily influenced by “leaders” that have much better designs than others (i.e., larger responses), while the diagonal elements of the leading agents are simultaneously decreased to maintain exploration. [24] provides such examples.

2) *Open Questions & Drawbacks*: The consensus framework raises many interesting questions: (i) Sharing $\mathbf{x}_{k,t}^c$ without its output may compromise privacy. While noise can be added, $\mathbf{x}_{k,t}^c + \epsilon_t$, what are the privacy-performance trade-offs? (ii) What if agents have varying fidelities? How can we incorporate this into $\mathbf{W}^{(t)}$ and perhaps learn the fidelities in the process? (iii) If resources vary, how can we restrict some agents' budgets to, say, $T - T'$ experiments while maintaining the effectiveness of consensus? (iv) Can we dynamically allow agents to borrow more strength from similar counterparts? Understanding heterogeneity for black-box functions is, of course, very challenging. Also, BO is designed for maximization, not for fully understanding the response surface in a way that provides a measure of heterogeneity. That said, can the observed improvements across iterations guide an agent on which other agent(s) they should depend on more?

a) *Note on theory*: We conclude by noting that, while consensus somewhat resembles predictive FL frameworks, deriving a theoretical foundation poses a critical open question (and perhaps a drawback). First, the black-box nature of f_k makes it difficult to derive theory. Second, despite recent advances in understanding the generalization error of GPs [27], understanding how these errors propagate to the, often non-concave, expected utility $F_{k,t}$ remains an open and challenging problem.

B. Conditioned Local Decisions

Global decisions dictate the next-to-observe designs to all agents, which may hinder collaboration in certain FL settings. Additionally, consensus requires all agents to collaborate actively and synchronously, which can be quite restrictive in many real-life scenarios. For example, consensus does not allow the use of good designs derived from historical data or human expertise.

Thus, an alternative framework for federated BO is to move decision-making locally, where an agent k makes local decisions conditioned on information $E_{k,t}$ from others agents

$k' \in [K \setminus k]$. Yet, what should $E_{k,t}$ be? The hope in FL is that $E_{k,t}$ is tiny in size but rich in information.

1) *Sharing Near-optimal Designs*: Following BO's philosophy to allocate designs near-optimal regions, and since we believe the black-box functions may have commonalities across agents, then they may share similar near-optimal regions. Therefore, if agents can help pinpoint each other to such regions, the problem simplifies then to finding individualized solutions within the region. With this rationale, one can define $E_{k,t}$ be a set of *synthetic* designs shared from other agents $k' \in [K \setminus k]$ to agent k , such that every borrowed design $\mathbf{x}_{k',t}^+ \in E_{k,t}$ satisfies the constraint $f_k(\mathbf{x}_{k',t}^+) > \delta_{k,t}$ for a constant $\delta_{k,t}$ picked by agent k .

Along this line, [28], chose $\mathbf{x}_{k',t}^+ \in E_{k,t}$ to be designs from other agents k' that satisfy the constraint

$$\mu_{k',t}(\mathbf{x}) - \eta_t \cdot \sigma_{k',t}(\mathbf{x}) > \delta_{k,t} \triangleq \max_{\mathbf{x}} \mu_{k,t}(\mathbf{x}), \quad (7)$$

where $\mu_{k,t}$ and $\sigma_{k,t}$ for an agent k are the mean and standard deviation of $\mathbb{P}_{\hat{f}_k | \mathcal{D}_{k,t}}$. Therefore, every design $\mathbf{x}_{k',t}^+ \in E_{k,t}$ points out a potential design with a better response compared to agent k 's current best, i.e., $\max_{\mathbf{x}} \mu_{k,t}(\mathbf{x})$, according to their posterior belief. In other words, any design in $E_{k,t}$ points out to a region for potential improvement. This is shown in Fig. 2. Critically, the comparison function in (7) can be efficiently and privately calculated for FL using secure multiparty computation, which is known as Yao's Millionaires' problem [29].

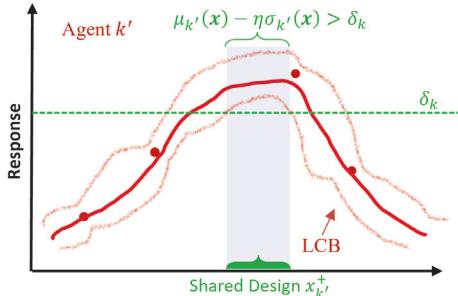


Fig. 2. Agent k' shares a design from the shaded region with agent k

A simple choice that satisfies (7), if it exists, is agent k' 's lower confidence bound (LCB) maximizer; $\mathbf{x}_{k',t}^+ = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \{\mu_{k',t}(\mathbf{x}) - \eta_t \cdot \sigma_{k',t}(\mathbf{x})\}$. It is worthwhile noting that designs (i) suggested by human experts or (ii) extracted from historical knowledge can simply be placed within $E_{k,t}$.

Now, agent k presumes that for all $\mathbf{x}_{k',t}^+ \in E_{k,t}$, there holds $f_k(\mathbf{x}_{k',t}^+) > \delta_{k,t}$. Notice that this constraint may not hold because of heterogeneity or \mathcal{GP} estimation error, but this can be readily accommodated as will become clear shortly. With $E_{k,t}$, agent k will update their posterior belief by conditioning on $E_{k,t}$ to make a local decision. The aim is to *fuse one's private data $\mathcal{D}_{k,t}$ with the shared knowledge $E_{k,t}$ that may be pointing to regions of improvement*.

Mathematically, the conditioned surrogate is given as $\mathcal{F}_{k,t}^+ \triangleq \mathbb{P}_{\hat{f}_k | \mathcal{D}_{k,t}} | E_{k,t}$. However, note that $E_{k,t}$ is not a dataset but a set of constraints, which essentially leads to a constrained \mathcal{GP} . Fortunately, sampling from a constrained \mathcal{GP} is well

studied, and one can employ basic rejection sampling (RS) to do so. A key benefit of RS in a constrained \mathcal{GP} is that one can set a finite sampling size, and if no samples are accepted at a borrowed design $\mathbf{x}_{k',t}^+$, the design is discarded. By doing so, agent k discards borrowed knowledge that significantly contradicts its posterior belief. This is the key to handling heterogeneity and demonstrating sublinear regret guarantees, regardless of how heterogeneous the functions are.

By extracting samples from $\mathcal{F}_{k,t}^+$, the expected utility $\mathbb{E}_{\mathcal{F}_{k,t}^+} [U(f(\mathbf{x}) | \mathcal{D}_{k,t})]$ can be approximated using Monte Carlo (see [28]), and a local decision $\mathbf{x}_{k,t}^{\text{new}}$ is made accordingly.

2) *Sharing Near-optimal Design Distributions*: While not currently proposed, an approach similar in spirit to sharing near-optimal designs could involve sharing distributional beliefs over the optimal design location. Mathematically, we can let $E_{k,t} = \{\pi_{k'}\}_{k'}$, where $\pi_{k'}$ is a density shared by agent k' . A simple example would be $\pi_{k'} = \mathcal{N}(\mathbf{x}_{k',t}^+, \sigma^2 \mathbf{I}_d)$ for some predetermined σ . Alternatively, agent k' can use repeated Thompson sampling to estimate their current belief about their optimal design, given as $\pi_{k'}(\mathbf{x}_k^*)$.

Now, instead of tweaking the surrogate as proposed in [28], one can directly adjust the expected utility $F_{k,t}$ in a manner similar to incorporating prior knowledge into BO in π BO [30]. With this, we can redefine the local objective to be:

$$\mathbf{x}_{k,t}^{\text{new}} = \operatorname{argmax}_{\mathbf{x}} \left[F_{k,t} \times \left(\prod_{k'} \pi_{k'} \right)^{\frac{\beta}{T}} \right] \quad (8)$$

The main idea here is to tweak the local objective $F_{k,t}$ using beliefs about the optimal design in $\pi_{k'}$. Given potential heterogeneity, the impact of information shared by other agents decays at a rate of β/T for some constant β . This allows an agent to give greater weight to designs that others consider optimal, while eventually prioritizing their own objective. Various adjustments to this approach are possible, such as taking a weighted product of the $\pi_{k'}$ values, with each $\pi_{k'}$ decaying at a different rate. Additionally, privacy can be preserved in a federated setting by adding additive or multiplicative noise to the shared density or its summary statistics.

With this approach, it is evident that updating the utility directly, rather than altering the surrogate, is mathematically simpler. However, it introduces the challenge of tuning β , which should be a function of the similarities across f_k and $f_{k'}$'s. Given the black-box nature of the functions, this can be a pathological problem. In contrast, the proposed approach of tweaking the surrogate $\mathbb{P}_{\hat{f}_k | \mathcal{D}_{k,t}} | E_{k,t}$ does not depend on tuning parameters and comes with strong guarantees simply since it enables discarding shared information that significantly conflict with an agent k 's own posterior belief.

3) *Sharing \mathcal{GP} Features*: Within the general framework of sharing and utilizing lightweight information $E_{k,t}$ for local decisions, [26] tackled federated BO by sharing \mathcal{GP} Random Fourier Features (RFFs) [31]. Using RFFs, stationary kernels ($\mathcal{K}(\mathbf{x}, \mathbf{x}') = \mathcal{K}(\mathbf{x} - \mathbf{x}')$) can be approximated by a set of random features $\phi(\mathbf{x}) \in \mathbb{R}^D$; where $\mathcal{K}(\mathbf{x}, \mathbf{x}') = \phi^\top(\mathbf{x}) \phi(\mathbf{x}')$. Therefore, the black-box function is approximated using

Bayesian linear regression as $\hat{f}_k(\mathbf{x}) = \phi^\top(\mathbf{x})\mathbf{w}_k$ such that the posterior of \mathbf{w}_k assuming a prior $\mathbf{w}_k \sim \mathcal{N}(0, \mathbf{I}_D)$ is

$$\mathbb{P}_{\mathbf{w}_k|\mathcal{D}_{k,t}} = \mathcal{N}(\boldsymbol{\nu}_{k,t}, \sigma^2 \boldsymbol{\Sigma}_{k,t}^{-1}),$$

where $\boldsymbol{\Phi}_{k,t} = [\phi(\mathbf{x}_{k,1}), \dots, \phi(\mathbf{x}_{k,t})]^\top$ is a $t \times D$ -dimensional matrix, $\boldsymbol{\Sigma}_{k,t} = \boldsymbol{\Phi}_{k,t}^\top \boldsymbol{\Phi}_{k,t} + \sigma^2 \mathbf{I}$, $\boldsymbol{\nu}_{k,t} = \boldsymbol{\Sigma}_{k,t}^{-1} \boldsymbol{\Phi}_{k,t}^\top \mathbf{y}_{k,[t]}$. Here σ^2 is the variance of ϵ and $\mathbf{y}_{k,[t]} = [y_{1,t}, \dots, y_{k,t}]^\top$.

With this, [26] proposed that $E_{k,t} = \{\mathbf{w}_{k'}^{(s)}; k' \in [K \setminus k]\}$, where $\mathbf{w}_{k'}^{(s)} \sim \mathbb{P}_{\mathbf{w}_{k'}|\mathcal{D}_{k',t}}$. This entails sharing the \mathcal{GP} coefficients, assuming ϕ is fixed across all agents. Now, agent k will either use their own $\mathbf{w}_k^{(s)}$ with probability p_t to find $\mathbf{x}_{k,t}^{\text{new}}$ using the Thompson sampling utility, i.e., $\mathbf{x}_{k,t}^{\text{new}} = \arg \max_{\mathbf{x}} \phi^\top(\mathbf{x})\mathbf{w}_k^{(s)}$, or use a $\mathbf{w}_{k'}^{(s)}$ randomly sampled from $E_{k,t}$ with probability $1 - p_t$ to get $\mathbf{x}_{k,t}^{\text{new}}$. Similar to the consensus idea, $p_t \rightarrow 1$ so that agents eventually focus on their own objectives. This idea was further expanded in [25] to guarantee differential privacy in FL by sharing an averaged coefficient $E_{k,t} = \{\bar{\mathbf{w}}_{k' \subset [K \setminus k]}\}$, which is clipped to bound its L_2 norm and has Gaussian noise added to it. The local process remained the same.

4) *Open Questions & Drawbacks:* While sharing optimal designs can incorporate any utility and provide theoretical guarantees despite the heterogeneity level, RS scales poorly with shared designs. Furthermore, sharing \mathcal{GP} RRFs is restricted to Thompson sampling and requires fixed features across all agents. Fortunately, the conditioned local design framework, where decisions are made at the agent level and informed by $E_{k,t}$, is *generic*, allowing for various ideas to be incorporated. Some open questions include: (i) When adding noise to $E_{k,t}$ (for RFFs and designs), what are the accuracy-privacy trade-offs? (ii) Can we include upper confidence bounds information when sharing designs to gauge which shared designs have the best potential? (iii) How can non-myopic BO literature be used to enhance collaboration? (iv) Can we theoretically identify informative conditions for when collaboration is helpful to develop algorithms that determine when to stop collaborating?

C. A Predictive Solution: Improving the Surrogate

The third framework does not tackle decision-making directly but instead aims to improve the local predictive model (i.e., the BO surrogate $\mathbb{P}_{\hat{f}_k|\mathcal{D}_{k,t}}$) through collaboration. Local decisions then follow the regular BO process as in (2). A central feature of this framework is that it can potentially be incorporated within the first two. For instance, consensus can be achieved using surrogates learned in a federated and collaborative manner. Here, two ideas arise:

1) *Federated \mathcal{GP} :* In BO, f_k is often modeled as a \mathcal{GP} where, $f_k(\mathbf{x}) \sim \mathcal{GP}(0, \mathcal{K}(\cdot, \cdot; \boldsymbol{\theta}_k))$ and $y_k = f_k(\mathbf{x}) + \epsilon_k$ where $\epsilon \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$. $\boldsymbol{\theta} = (\boldsymbol{\theta}_k, \sigma^2)$ parameterize the \mathcal{GP} and encode our prior belief. Naturally, $\boldsymbol{\theta}$ could be learned in a federated fashion. Here the utility function $L_{k,t}$ for each agent is the log-marginal likelihood (9). Now, in FL, the goal is to collaboratively learn $\boldsymbol{\theta}$ that maximizes the global utility

$$\begin{aligned} \underset{\boldsymbol{\theta}}{\operatorname{argmax}} \sum_{k=1}^K p_k L_{k,t} &\triangleq \\ \sum_{k=1}^K p_k \log \mathbb{P}(\mathbf{y}_{k,[t]}|\mathcal{D}_{k,t}; \boldsymbol{\theta}) . \end{aligned} \quad (9)$$

Recently, [32] proposed FedAvg to solve (9). Unlike in deep learning, $L_{k,t}$ is a log-likelihood function featuring *correlations* (due to the \mathcal{GP}), which leads to stochastic gradients being biased estimators of the full gradient. Nevertheless, [11] shows that despite these correlations, FedAvg converges to a critical point of (9), subject to statistical errors. A key advantage of a federated \mathcal{GP} is its inherent ability to personalize, as estimating $\boldsymbol{\theta}$ is equivalent to learning a global \mathcal{GP} prior. Personalized predictions $\mathbb{P}(y_k^*|\mathcal{D}_{k,t}, \mathbf{x}^*; \boldsymbol{\theta})$ are obtained by conditioning on a local agent's data $\mathcal{D}_{k,t}$. The local data inherently fulfill the role of personalization.

2) *Federated Multi-output \mathcal{GP} (\mathcal{MGP})*: Alternatively, one may give more model flexibility by assuming agents have shared $\boldsymbol{\theta}^g$ and unique parameters $\boldsymbol{\theta}^k$ through an \mathcal{MGP} where

$$[f_1(\mathbf{x}), \dots, f_K(\mathbf{x})]^\top \sim \mathcal{MGP}(\mathbf{0}, \text{cov}_{k,k'}^f(\mathbf{x}, \mathbf{x}')) . \quad (10)$$

Many options can be used to construct $\text{cov}_{k,k'}^f(\mathbf{x}, \mathbf{x}')$. Perhaps, the most common approach is convolution processes [33] $f_k(\mathbf{x}) = \sum_{q=1}^Q \int_{-\infty}^{\infty} \mathcal{K}_{q,k}(\mathbf{x} - \mathbf{u}; \boldsymbol{\theta}^k) g_q(\mathbf{u}; \boldsymbol{\theta}^g) d\mathbf{u}$ where g is a latent \mathcal{GP} . The key advantages is that multiple latent functions g_q allow information sharing across different agents through different kernels $\mathcal{K}_{q,k}(\mathbf{x})$.

Unlike federated \mathcal{GP} s, the global \mathcal{MGP} utility cannot be written as a sum of agent utilities like (9), since agents are all correlated by construction. Despite this, [34] recently show that, using variational inference, one can approximate the joint \mathcal{MGP} log-likelihood as $\sum_{k=1}^K p_k \tilde{L}_{k,t}$, where $\tilde{L}_{k,t}$'s are independent. With this, they developed a simple FedAvg-like approach to update $\{\{\boldsymbol{\theta}^k\}, \sigma^2, \boldsymbol{\theta}^g\}$ simultaneously, while only sharing $\boldsymbol{\theta}^g$ with the orchestrator.

3) *Key Open Question*: A central question for integratively learning surrogates is: What is the best way to initialize federated BO? Should agents begin with space-filling designs, or should they explore different space partitions? Furthermore, how will this affect \mathcal{GP} estimation error early on?

IV. CONCLUSION

We present a vision and unifying frameworks for collaborative and federated black-box optimization from a Bayesian optimization perspective while highlighting the many exciting problems yet to be addressed. We conclude by noting that many applications could benefit from this emerging literature, including collaborative optimization of manufacturing process parameters, material discovery, hyperparameter tuning and clinical trials. Therefore, we believe a high-impact endeavor would be to bring this theory to practical applications by providing a working example where agents collaborate to improve their decisions and benefit from collaboration, all while preserving their privacy and intellectual property.

REFERENCES

[1] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh, "Scaffold: Stochastic controlled averaging for federated learning," in *ICML*. PMLR, 2020, pp. 5132–5143.

[2] H. T. Nguyen, V. Sehwag, S. Hosseinalipour, C. G. Brinton, M. Chiang, and H. V. Poor, "Fast-convergent federated learning," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 1, pp. 201–218, 2020.

[3] F. Sattler, S. Wiedemann, K.-R. Müller, and W. Samek, "Robust and communication-efficient federated learning from non-iid data," *IEEE transactions on neural networks and learning systems*, vol. 31, no. 9, pp. 3400–3413, 2019.

[4] D. Li and J. Wang, "Fedmd: Heterogenous federated learning via model distillation," *NeurIPS 2019 Workshop*, 2019.

[5] H. Wang, M. Yurochkin, Y. Sun, D. Papailiopoulos, and Y. Khazaeni, "Federated learning with matched averaging," *ICLR*, 2020.

[6] N. Shi and R. A. Kontar, "Personalized federated learning via domain adaptation with an application to distributed 3d printing," *Technometrics*, vol. 65, no. 3, pp. 328–339, 2023.

[7] A. Fallah, A. Mokhtari, and A. Ozdaglar, "Personalized federated learning: A meta-learning approach," *NeurIPS 33*, 2020.

[8] A. N. Bhagoji, S. Chakraborty, P. Mittal, and S. Calo, "Analyzing federated learning through an adversarial lens," in *ICML*. PMLR, 2019, pp. 634–643.

[9] T. Li, M. Sanjabi, A. Beirami, and V. Smith, "Fair resource allocation in federated learning," *ICLR*, 2019.

[10] X. Yue, M. Nouiehed, and R. Al Kontar, "Gifair-fl: A framework for group and individual fairness in federated learning," *INFORMS Journal on Data Science*, vol. 2, no. 1, pp. 10–23, 2023.

[11] X. Yue, R. A. Kontar, and A. M. E. Gómez, "Federated data analytics: A study on linear models," *IIE Transactions*, vol. 56, no. 1, pp. 16–28, 2024.

[12] N. Shi, F. Lai, R. Al Kontar, and M. Chowdhury, "Fed-ensemble: Ensemble models in federated learning for improved generalization and uncertainty quantification," *IEEE Transactions on Automation Science and Engineering*, 2023.

[13] R. Kontar, N. Shi, X. Yue, S. Chung, E. Byon, M. Chowdhury, J. Jin, W. Kontar, N. Masoud, and M. Nouiehed, "The internet of federated things (ioft)," *IEEE Access*, vol. 9, pp. 156 071–156 113, 2021.

[14] N. Shi and R. A. Kontar, "Personalized pca: Decoupling shared and unique features," *Journal of Machine Learning Research*, vol. 25, no. 41, pp. 1–82, 2024. [Online]. Available: <http://jmlr.org/papers/v25/22-0810.html>

[15] N. Shi, R. Al Kontar, and S. Fattah, "Heterogeneous matrix factorization: When features differ by datasets," *arxiv:2305.17744*, 2023.

[16] N. Shi, S. Fattah, and R. A. Kontar, "Triple component matrix factorization: Untangling global, local, and noisy components," *Journal of Machine Learning Research*, vol. 25, no. 332, pp. 1–76, 2024. [Online]. Available: <http://jmlr.org/papers/v25/24-0400.html>

[17] G. Liang, N. Shi, R. Al Kontar, and S. Fattah, "Personalized dictionary learning for heterogeneous datasets," *NeurIPS*, vol. 36, 2024.

[18] J. González, Z. Dai, P. Hennig, and N. Lawrence, "Batch bayesian optimization via local penalization," in *AISTATS*. PMLR, 2016, pp. 648–657.

[19] J. Azimi, A. Fern, and X. Fern, "Batch bayesian optimization via simulation matching," *NeurIPS*, vol. 23, 2010.

[20] P. I. Frazier, "A tutorial on bayesian optimization," *arXiv preprint arXiv:1807.02811*, 2018.

[21] R. B. Gramacy and H. K. Lee, "Adaptive design and analysis of supercomputer experiments," *Technometrics*, vol. 51, no. 2, pp. 130–145, 2009.

[22] D. R. Jones, M. Schonlau, and W. J. Welch, "Efficient global optimization of expensive black-box functions," *Journal of Global optimization*, vol. 13, no. 4, p. 455, 1998.

[23] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *AISTATS*. PMLR, 2017, pp. 1273–1282.

[24] X. Yue, R. A. Kontar, A. S. Berahas, Y. Liu, Z. Zai, K. Edgar, and B. N. Johnson, "Collaborative and distributed bayesian optimization via consensus: Showcasing the power of collaboration for optimal design," *arXiv preprint arXiv:2306.14348*, 2023.

[25] Z. Dai, B. K. H. Low, and P. Jaillet, "Differentially private federated bayesian optimization with distributed exploration," *NeurIPS*, vol. 34, pp. 9125–9139, 2021.

[26] Z. Dai, K. H. Low, and P. Jaillet, "Federated bayesian optimization via thompson sampling," *NeurIPS*, 2020.

[27] A. Lederer, J. Umlauf, and S. Hirche, "Uniform error bounds for gaussian process regression with application to safe control," *NeurIPS*, vol. 32, 2019.

[28] Q. Chen, L. Jiang, H. Qin, and R. A. Kontar, "Multi-agent collaborative bayesian optimization via constrained gaussian processes," *Technometrics*, no. just-accepted, pp. 1–23, 2024.

[29] H.-Y. Lin and W.-G. Tzeng, "An efficient solution to the millionaires' problem based on homomorphic encryption," in *ACNS*, vol. 5. Springer, 2005, pp. 456–466.

[30] C. Hvarfner, D. Stoll, A. Souza, M. Lindauer, F. Hutter, and L. Nardi, "pibo: Augmenting acquisition functions with user beliefs for bayesian optimization," *arXiv preprint arXiv:2204.11051*, 2022.

[31] A. Rahimi and B. Recht, "Random features for large-scale kernel machines," *NeurIPS*, vol. 20, 2007.

[32] X. Yue and R. Kontar, "Federated gaussian process: Convergence, automatic personalization and multi-fidelity modeling," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.

[33] M. A. Alvarez and N. D. Lawrence, "Computationally efficient convolved multiple output gaussian processes," *The Journal of Machine Learning Research*, vol. 12, pp. 1459–1500, 2011.

[34] S. Chung and R. Al Kontar, "Federated multi-output gaussian processes," *Technometrics*, vol. 66, no. 1, pp. 90–103, 2024.