

Efficient Visual Appearance Optimization by Learning from Prior Preferences

Zhipeng Li
Department of Computer Science
ETH Zürich
Zürich, Switzerland
zhipeng.li@inf.ethz.ch

Yi-Chi Liao
Department of Computer Science
ETH Zürich
Zürich, Switzerland
yichi.liao@inf.ethz.ch

Christian Holz
Department of Computer Science
ETH Zürich
Zürich, Switzerland
christian.holz@inf.ethz.ch




Figure 1: We introduce *Meta-PO*, a computational method that models users’ implicit preferences of visual appearances given a theme (here “cold”) as they iteratively select from a sampled set of candidates presented in a 2D gallery. *Meta-PO* integrates Preferential Bayesian Optimization with meta-learning to increase optimization efficiency by leveraging population models derived from *prior optimization experiences* — across users and themes. *Meta-PO* can thus produce more satisfying results in fewer iterations.

ABSTRACT

Adjusting visual parameters such as brightness and contrast is common in our everyday experiences. Finding the optimal parameter setting is challenging due to the large search space and the lack of an explicit objective function, leaving users to rely solely on their implicit preferences. Prior work has explored Preferential Bayesian Optimization (PBO) to address this challenge, involving users to iteratively select preferred designs from candidate sets. However, PBO often requires many rounds of preference comparisons, making it more suitable for designers than everyday end-users. We propose *Meta-PO*, a novel method that integrates PBO with meta-learning to improve sample efficiency. Specifically, *Meta-PO* infers prior users’

preferences and stores them as models, which are leveraged to intelligently suggest design candidates for the new users, enabling faster convergence and more personalized results. An experimental evaluation of our method for appearance design tasks on 2D and 3D content showed that participants achieved satisfactory appearance in 5.86 iterations using *Meta-PO* when participants shared similar goals with a population (e.g., tuning for a “warm” look) and in 8 iterations even generalizes across divergent goals (e.g., from “vintage”, “warm”, to “holiday”). *Meta-PO* makes personalized visual optimization more applicable to end-users through a generalizable, more efficient optimization conditioned on preferences, with the potential to scale interface personalization more broadly.



This work is licensed under a Creative Commons Attribution 4.0 International License.
UIST '25, September 28–October 1, 2025, Busan, Republic of Korea
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2037-6/2025/09
<https://doi.org/10.1145/3746059.3747622>

CCS CONCEPTS

• Computing methodologies → Machine learning; Graphics systems and interfaces; • Human-centered computing;

KEYWORDS

Human-in-the-loop optimization, visual design, preferential optimization, preferential modeling, Bayesian optimization, meta-learning, meta-Bayesian optimization

ACM Reference Format:

Zhipeng Li, Yi-Chi Liao, and Christian Holz. 2025. Efficient Visual Appearance Optimization by Learning from Prior Preferences. In *The 38th Annual ACM Symposium on User Interface Software and Technology (UIST '25)*, September 28–October 1, 2025, Busan, Republic of Korea. ACM, New York, NY, USA, 24 pages. <https://doi.org/10.1145/3746059.3747622>

1 INTRODUCTION

Adjusting visual parameters is a common part of everyday life. Most people tune image appearance by adjusting brightness, contrast, or applying filters, before sharing photos on social media or for professional use [7]. However, this task is challenging for two reasons. First, the search space is overwhelmingly large. There are 10 to 15 common parameters that affect image appearance, e.g., exposure, saturation, hue, and sharpness. Even with just 10 parameters discretized into 10 levels each, the number of combinations reaches 10^{10} , making a thorough manual search impractical. Second, visual design lacks an explicit objective function. The “goodness” of a design is subjective, governed by the user’s implicit preferences. Thus, end-users rarely try to find their ideal appearance from scratch, and instead rely on pre-defined filters [42]. Even professional designers, though able to spend more time exploring, often settle for a result that is “good enough” without a clear sense of optimality [11, 30].

To address this challenge, prior works introduced Preferential Bayesian Optimization (PBO) as a principal approach, along with applying existing filters and tuning parameters. PBO builds on Bayesian Optimization (BO), a probabilistic framework for solving black-box functions whose internal structure is inaccessible and must be learned through iterative evaluations [17, 45]. BO is widely used for human-in-the-loop tasks due to its robustness, generalizability, and ability to handle uncertainty [28, 53]. PBO extends BO to settings where the values of the objective functions are *not directly observable*, relying instead on inferred user preferences through comparative feedback. Practically, PBO operates via selection-based interaction: users choose the most preferred design among a presented set. By observing these selections, the optimizer gradually learns a model of the user’s implicit preferences, eventually converging on the most preferred design. With this mechanism, PBO enables users to navigate design spaces where manual tuning is impractical due to high dimensionality or a lack of pre-defined filters. This capability has demonstrated its effectiveness and efficiency in various domains, including fabrication [13], font design [57], and color selection [22].

Prior work has proposed various interactions to improve sample efficiency in image enhancement using PBO. For example, linear search spaces [27] and gallery-based interfaces [26] have been used to elicit preferences beyond the typical pairwise comparisons [19]. Among these, the gallery-based interaction (illustrated in Figure 1) achieves the highest efficiency, as it enables exploration of a 2D search plane in one iteration. These studies also showed that PBO allows users to produce images that are more preferred than those

created using pre-defined filters in Photoshop and Lightroom. Despite these advances, PBO still requires 10 to 15 iterations to converge [26], potentially limiting its practicality in everyday applications. In this paper, we aim to enable PBO to *converge even more efficiently toward a user’s implicit preference* by leveraging prior users’ preference data, making it suitable for end-users in everyday visual design tasks. This idea is rooted in transfer learning (adapting to new but related tasks) and meta-learning (learning to learn). While both have been widely studied in machine learning and optimization [2, 60, 64], only recently have they been explored in human-in-the-loop tasks [36]. To our knowledge, no prior work has applied meta-learning to preferential optimization, where preferences are inferred via comparisons.

We propose a novel computational method, Meta-PO (short for “**Meta**-learning for **P**referential **B**ayesian **O**ptimization”), to enable efficient visual appearance optimization. The core principle of Meta-PO is to transfer prior users’ preference models (learned from past optimization processes) on a specific theme (e.g., mood, theme, or filter) to accelerate the optimization process for new users with a similar target theme. For example, if prior users have previously completed PBO sessions tuning images to offer a “warm” mood, we construct separate models from their data. These models can then be leveraged to boost optimization efficiency for a new user aiming to achieve a similar “warm” visual effect. Our assumption is that while individual interpretations of “warm” vary slightly and target images differ, there are shared **across-users** preference patterns, which can support the future optimization. Furthermore, we extend this concept to a more challenging setting: **cross-theme generalization**. Suppose we have accumulated preference optimization data across diverse moods, themes, or filters (e.g., warm, winter, summer, happy, sad, vintage, autumn, and more). When a new user aims to adjust images for an unseen goal (e.g., “cold”), our method can still utilize relevant prior data, as shown in Figure 1. Meta-PO achieves this by automatically identifying the relevance of prior models to adjust their influence level on the current optimization, allowing knowledge transfer from related themes while minimizing interference from unrelated ones. Moreover, Meta-PO has an internal mechanism where the guidance of prior models is stronger in the initial phase of an optimization, and gradually, the outcome will be tailored only based on the observations from the new user.

Several computational elements are integrated to build Meta-PO. Our interaction is built on the Sequential-Gallery [26], where users are presented with a 2D plane of visual design options sampled from a high-dimensional design space, allowing them to select the most preferred design (the interaction is illustrated in Figure 1). We incorporate an aggregation-based meta-learning mechanism: a set of prior preferential models is constructed from previously collected user data, and these models collectively guide the ongoing optimization. Each model’s influence is determined by dynamically assigned weights, with higher weights indicating stronger influence. Specifically, we adopt a variant of the Transfer Acquisition Function (TAF-R) [67], which computes similarity between the current user’s preferences and those of prior users, enabling selective and adaptive knowledge transfer. To construct the 2D gallery [26], we require at least three design points. As in the original method, we select the current best design (the user’s most preferred so far) as the first point. We then compute the second point using TAF-R,

leveraging meta-learned knowledge to guide the search. Our final novelty lies in the selection of the third design point. We aim to, again, leverage prior knowledge without proposing a point too similar to the second. We combine TAF-R with the 2-step Bayesian optimization method [68], which identifies a candidate likely to provide significant improvement if the current most promising design were to be evaluated. This ensures diversity and exploration in the 2D gallery presented to the user.

We evaluate Meta-PO’s **cross-user transfer** and **cross-theme generalization** through user studies, demonstrating improved efficiency on two distinct applications: image enhancement and virtual reality lighting. With prior optimization experience on the same theme, participants adjusted images to their desired visual theme in 5.86 ($SD = 1.20$) iterations, compared to 9.54 ($SD = 2.19$) without leveraging previous experience. In cross-theme scenarios, leveraging optimization experience from different themes, they required 7.41 ($SD = 1.28$) iterations, also significantly fewer than when no prior experience was utilized. Participants also reported that Meta-PO enables more effective exploration and more satisfying outcomes. In virtual lighting design, Meta-PO users achieved satisfactory results in fewer iterations ($M = 4.05$, $SD = 1.39$) than those without knowledge transfer ($M = 6.85$, $SD = 1.48$), demonstrating its effectiveness in this application as well. To summarize, we make the following contributions:

- We introduce Meta-PO, a novel method incorporating several computational elements to enable efficient preference-driven optimization, allowing both end-users and designers to identify their most preferred designs within just a few selections.
- We demonstrate the effectiveness of Meta-PO in achieving cross-user and cross-theme adaptation across both 2D and 3D visual design tasks, showcasing its generalizability and efficiency in real-world scenarios.

2 RELATED WORK

2.1 Preferential Bayesian Optimization

Optimizing human-centered interfaces is challenging because the only information we have regarding the quality of the design is acquired via user evaluation. Bayesian optimization (BO) has long been employed to support this process. Formally, the optimization process seeks $x^* = \arg \max_{x \in \mathcal{X}} f(x)$, where x is the parameter vector, and $f(\cdot)$ is the human related objective function. In each iteration, the optimizer suggests a parameter setting, and the user provides corresponding feedback. BO relies on two fundamental elements: (1) a *surrogate model*, typically a Gaussian Process (GP), which approximates the objective function, and (2) an *acquisition function*, which selects the next evaluation point. These two components enable BO to effectively manage the uncertainty in human-related objective functions, achieving a balance between exploration and exploitation under limited evaluation budgets. This capability makes BO particularly well-suited for human-in-the-loop tasks, where evaluations involving human feedback tend to be expensive and time-consuming [6, 37, 65]. Leveraging these strengths, BO has found applications in a wide range of tasks because of its robust performance [17, 53, 63], including input techniques [36], designing software interfaces [44] and physical interactions [37, 38],

refining haptic interfaces [5], game design [25], UI design in urban air mobility [43], and other human-involved interactions.

In the above examples, BO directly observes the measurable quality of the design. However, there is another type of task that is unrealistic to expect human users to provide precise quantitative assessments. For example, asking users to tell the goodness of a parameter configuration for adjusting photos’ presentation. The goodness function is implicit in the human mind, which can not be measured directly. Prior works addressed this challenge by Preferential Bayesian Optimization (PBO) [19]. Instead of requiring numerical evaluations, PBO asks users to compare the results of two configurations (x_a and x_b) and express a preference (e.g., $f(x_a) > f(x_b)$ or $f(x_a) < f(x_b)$). These comparisons are then modeled probabilistically using approaches such as the Bradley-Terry-Luce (BTL) model [27, 59], allowing PBO to integrate preference-based feedback within the standard BO framework to suggest the next evaluation candidates. This makes PBO achieve greater efficiency and robustness compared to standard BO that relies on direct numerical ratings, especially when optimizing subjective or implicit functions [19, 56]. Prior works explored interactive approaches that allow the users to identify the most preferred design among a set of options to further improve the efficiency [27, 70], which will be reviewed in Preliminaries.

Although PBO is effective, its initial exploration phase is inefficient due to the lack of prior knowledge, requiring extensive random sampling [63]. However, visual appearance functions often exhibit similarities across tasks. Leveraging this structure, we propose using meta-learning to accelerate early exploration and improve efficiency in visual appearance design.

2.2 Meta-learning for Bayesian optimization

Meta-learning is a machine learning paradigm that aims to enable models to quickly adapt to new tasks by leveraging prior experience from similar tasks [52, 62]. It has effectively improved the adaptation efficiency for recognition tasks and reinforcement learning [15, 16, 46]. Recent research in BO has attempted to integrate the concept of meta-learning into black-box optimization problems. That is, improving optimization efficiency by incorporating knowledge from previous optimization processes for similar problems. While Meta-Bayesian optimization (Meta-BO) has been employed in hyperparameter tuning and various machine learning tasks [32, 33, 63], its application to human-computer interaction (HCI) problems has only recently been explored [36].

There are multiple approaches to realizing Meta-BO. A straightforward method is to integrate data from the current task into a surrogate model (typically a Gaussian Process) that has been trained on previous tasks, which provides strong prior knowledge [3, 55, 71]. However, GPs do not scale well to large datasets due to their cubic computational complexity; that is, $O(n^3)$, where n represents the total number of observations across all tasks. Another approach involves replacing the surrogate model or acquisition function with deep neural networks to handle large datasets [20, 61, 66]. However, these methods typically require substantial amounts of training data and generally lack explainability, which is challenging or even impractical when the data requires human evaluations. Recently, Liao et al. proposed CHiLO, which uses a BNN as a population

model trained on samples from individual users' GP, reducing the need for large amounts of user data while preserving data quality [40]. However, directly applying this approach to preferential tasks may be problematic, as it assumes all users' functions are inherently similar and can be captured by a single model. In reality, human preferences often vary widely across users and tasks, and this might enforce the shared model risks over-smoothing user-specific differences.

Another approach is the weighted-sum method, which constructs individual GP models for each task and aggregates their acquisition functions based on weighted contributions [33, 67]. This approach is more practical compared to the other methods: it does not introduce cubically increasing complexity or assume a large amount of training data. Liao et al. showed this method can efficiently adapt the input device to replace manual calibration [36]. However, these methods have focused on optimizing measurable objectives, such as speed or accuracy. Extending it to preference-based tasks remains an open challenge for efficiently optimizing preference objectives. In this work, we took this approach but incorporated novel extensions for addressing the specific needs of preference-driven tasks. PBO must account for individual user preferences, which can vary significantly and even conflict at times. Further, the number of suggested designs increases in the PBO, requiring additional mechanisms. We address this by using multi-step BO, which will be detailed in Preliminaries.

2.3 Model-based Optimization for Design

Design is a core topic in HCI, with computational methods increasingly used to support designers and users in creating effective, user-friendly systems [21]. A widely adopted approach is model-based optimization, which formulates design problems as mathematical optimization tasks and evaluates candidates using predictive models rather than direct human feedback [35]. Unlike human-in-the-loop optimization methods, model-based approaches rely on surrogate models to assess design quality, enabling faster iteration and scalable design generation. These methods have been successfully applied in domains where reliable physical models are available to guide design decisions. For example, they have been used to optimize geometry for acoustic control [31], rotational inertia [1], and object balance [49], while respecting aesthetic and structural constraints. In UI design, optimization has improved keyboard layouts [24] and enabled real-time layout refinement through predictive models [47, 58]. In Mixed Reality (MR), where interfaces extend beyond traditional screens, optimization techniques help adapt UI layouts to real-world spatial contexts [8, 34]. Despite their broad applicability, model-based methods face two key limitations. First, they require task-specific models, which may be infeasible or unavailable for many human-centered tasks. Second, user preferences are inherently subjective and vary widely across individuals. While aesthetic models exist to capture broad stylistic attributes such as vibrancy, warmth, or minimalism [12, 14], they fall short in modeling individual differences of preferences, ultimately limiting true personalization.

3 PRELIMINARIES

We now review the areas that laid the foundation for Meta-PO. First, we review the development and underlying method of Preferential Bayesian Optimization (PBO). We then review the interactive systems that extend PBO for interactive systems, particularly focusing on the Sequential Gallery Interaction [26]. Furthermore, we review a specific meta-BO method, Transfer Acquisition Function (TAF) [67], which is the underlying approach for our method. Finally, we review a practical approach for 2-step Bayesian optimization [68] that is a critical part of Meta-PO to construct a user search plane at each iteration.

3.1 Modeling Human's Implicit Preference

The goal of PBO is to model implicit human preferences and optimize design accordingly. There are two approaches to eliciting these preferences. A standard and more robust method is to present users with a set of candidates and ask them to choose their preferred option. This comparative format is generally easier for users to engage with and provides richer, more informative feedback, from which PBO can infer preference models. Prior studies have shown this approach to be effective for preference-based optimization tasks [26, 27].

A simpler alternative is to ask users to provide scalar ratings indicating how much they like a given option. However, this method suffers from fundamental limitations. First, subjective ratings are often noisy and inconsistent: it is inherently difficult for users to translate their preferences into precise numerical values, and individuals tend to apply different rating scales. Moreover, as identified in prior work [39], scalar ratings are strongly biased by recent experiences, further reducing reliability. Second, scalar-based feedback is less efficient. In line-based or gallery-based PBO, each iteration yields multiple pairwise comparisons, accelerating convergence. In contrast, scalar ratings produce only a single data point per iteration, offering less information and requiring significantly more rounds of interaction. For these reasons, scalar-based PBO is rarely used in practice, except in settings where displaying multiple design candidates is infeasible, such as in driving scenarios [10].

We adopt the typical PBO approach by presenting a set of candidates to users, who compare them and select their most preferred choice. The selection results can be mathematically modeled using stochastic models such as the Bradley-Terry-Luce (BTL) model [59]. With the probabilistic representation of the selection results, PBO can fit a Gaussian Process (GP) with the standard Bayesian Optimization (BO) framework and decide the next set of candidate options with its acquisition function for the subsequent iterations. As more selection results are accumulated over time, the Gaussian Process progressively refines its model of the user's preferences.

The core difference between PBO and standard BO is that PBO has to translate the selection results into likelihoods. The BTL model, which has been widely used in previous PBO methods, is well-suited for handling such selection scenarios [59]. Specifically, consider m design candidates presented to the user, denoted as $\mathcal{P} = \{x_i\}_{i=1}^m$, from which the user selects x^c based on his preference. Then we collect the feedback in the following form:

$$x^c \succ \mathcal{P} \setminus \{x^c\} \quad (1)$$

where $\succ$ indicates that the perceived goodness value of x^c is greater than that of any other candidates in $\mathcal{P}$. The BTL model defines the likelihood of this outcome as:

$$p(x_i \succ \mathcal{P} \setminus x_i \mid \{g_j\}_{j=1}^m) = \frac{\exp(g_i)}{\sum_{j=1}^m \exp(g_j)} \quad (2)$$

where g_j represents the users' implicit goodness function for design candidate j . This allows us to fit the hyperparameters θ of the Gaussian Process by maximizing the posterior distribution as such:

$$\theta = \arg \max_{\theta} p(\theta \mid \mathcal{D}_k) = \arg \max_{\theta} \frac{p(\mathcal{D}_k \mid \theta) p(\theta)}{p(\mathcal{D}_k)}. \quad (3)$$

Here, $\mathcal{D}_k$ comprises all the prior user's preference responses until k -th iteration, and $p(\mathcal{D}_k)$ is a normalizing constant. $p(\theta)$ is the prior distribution over θ , which is predefined, and $p(\mathcal{D}_k \mid \theta)$ is the posterior distribution after all the observations, which can be rewritten as

$$p(\mathcal{D}_k \mid \theta) = \int p(\mathcal{D}_k \mid g) p(g \mid \theta) dg \quad (4)$$

where $p(g \mid \theta)$ is the prior probability defined by θ . Since $\mathcal{D}_k$ consists of a series of selection results, we compute the likelihood of the full dataset as a product over all comparisons observed so far. Let each iteration i involve a user selecting the preferred design x_i^c from a candidate set $\mathcal{P}_i$. Then:

$$p(\mathcal{D}_k \mid g) = \prod_{i=1}^k p(x_i^c \succ \mathcal{P}_i \setminus \{x_i^c\} \mid g) \quad (5)$$

With this formulation, the hyperparameters θ can be estimated using the series of preference-based selection results observed up to iteration k . A detailed procedure for fitting the Gaussian Process using the BTL model can be found in [27].

3.2 Interactions with PBO

With the implicit preference model, we still need to decide the next set of candidates and present them to users for selection. Prior works have implemented interactive systems that allow users to easily explore a low-dimensional search space constructed from the preference model and select the most preferred design instance. Sequential-Line adopts a linear search space, while Sequential-Gallery presents a 2D gallery of design candidates [26, 27]. The gallery-based interaction interface has been shown to encourage more effective exploration and inspire users to try diverse options, resulting in a more engaging and creative experience compared to traditional filter-based methods [26].

Sequential-Line uses two design points in the high-dimensional parameter space to define the user's search space at each iteration. One point is the design with the highest acquisition value, x^{AF} , computed using the GP model based on the user's selections so far. The second point is the best candidate observed to date, x^+ . Connecting these two points forms a line that defines the search space. Users can explore along this line using a slider and select their most preferred design. However, searching a high-dimensional space with a single slider typically requires many iterations. Sequential-Line addresses this by limiting the number of parameters to six and relying on crowd-sourced data in their user study [27].

To further boost the efficiency, Sequential-Gallery was proposed later, which presents candidates in the form of a 2D gallery. With the gallery arranged on a 2D plane, users can explore two dimensions in each iteration. To define the 2D search plane, Sequential-Gallery uses x^+ as the center, x_1^{AF} as one of the corners, and then searches for another point x_2^{AF} as another corner, which has the highest acquisition value in the orthogonal direction to the line connecting x^+ and x_1^{AF} . Additionally, Sequential-Gallery defines a symmetric search space, so the 2D gallery is defined by these three points. Sequential-Gallery then evenly samples discrete points in this space and presents them to the users for selection.

In this work, we follow the interaction of Sequential-Gallery, extracting a 2D search space from a high-dimensional design space. However, instead of relying solely on the current GP to construct the search plane, we focus on leveraging prior experiences as well as a two-step BO approach to suggest design points that can boost the sample efficiency, which will be expanded in the next sub-sections.

3.3 Transfer Acquisition Function (TAF): a Meta-Learning Method

We aim to enhance PBO's performance and efficiency via meta-learning. Specifically, we build on the Transfer Acquisition Function (TAF) [67], which maintains a separate model for each prior task, which is referred to as *population models*. During deployment, TAF also constructs a *current model*, a GP trained only on observations from the ongoing task. Here, a task refers to a user completing a preferential optimization session on an image or 3D scene with a specific target theme. To guide optimization, TAF aggregates acquisition values from both the population and current models using a weighted-sum approach. The candidate with the highest aggregated score is selected for evaluation.

We follow the prior works to use Expected Improvement (EI) as the basic acquisition function [36, 67]. EI computes the expected improvement of all candidates over the best observed objective value in the previous observations $\mathcal{D}_k$, as defined by

$$EI(x, \mathcal{D}_k) = E[I(x, \mathcal{D}_k)] = E[\max(0, \mu(x) - y_{\max})] \quad (6)$$

where $I(x)$ represents the predicted improvement, defined as the difference between the predicted mean $\mu(x)$ at a sampling point x and the best objective value observed $y_{\max}$. By definition, EI is the expected value of $I(x)$, indicating the anticipated improvement over the current best observation. Then, the TAF with EI as the base acquisition function could be computed with:

$$a(x) = \frac{w_{M+1} E[I_{M+1}(x)] + \sum_{i=1}^M w_i I_i(x)}{\sum_{i=1}^{M+1} w_i} \quad (7)$$

where M denotes the number of prior models, $I_i(x)$ is the improvement at sampling point x , as defined in Equation 6, and w_i represents the weight assigned to the i -th prior model.

Weights determine each model's influence: higher weights yield greater impact. Liao et al. [36] employed a technique, TAF-M, that assigns weights based on prediction confidence: higher confidence yields higher weight. However, this approach does not adapt to the current user's feedback or ongoing observations. We instead adopt TAF-R, a variant that determines weights based on similarity

between the current model and each population model. Details are provided in subsection 4.3.

3.4 Two-step Acquisition Function

To construct a 2D search plane for sequential gallery-based optimization, we require three design points. We compute one of these points by the two-step lookahead acquisition function [18, 29, 48, 68]. Here, we provide a short introduction to this advanced technique in BO. Standard BO uses a one-step greedy strategy, selecting the next point based solely on the acquisition function given the current model. In contrast, two-step BO anticipates how the next evaluation might reshape the model and impact subsequent choices. It does this by simulating the outcome of a potential first evaluation, updating the model hypothetically, and computing the second-step acquisition value conditioned on this simulated update.

Theoretically, this approach explicitly searches for the most effective decisions in the next round. Rather than merely exploiting current knowledge, two-step BO balances exploration and exploitation more effectively, making it less prone to local optima. Empirically, Wu and Frazier [68] demonstrates that two-step BO consistently outperforms one-step strategies across various optimization tasks. In our specific case, since the one design point is already selected by the typical one-step greedy acquisition strategy, employing a two-step lookahead acquisition for another point complements it by encouraging more principled and forward-looking exploration.

We build on Wu and Frazier [68], using Monte Carlo sampling to simulate the potential outcome of evaluating a first candidate point, then updating the Gaussian Process model with this simulated result. This is more practical as it efficiently approximates the first evaluation using posterior samples, avoiding expensive exact inference. Using Expected Improvement (EI) as an example, the two-step acquisition function can be written as:

$$a^{2\text{-step}}(x) = EI_0(X_1) + E \left[\max_{x_2 \in X} EI_1(x_2) \right] \quad (8)$$

X_1 denotes the simulated candidate(s) in the first step, EI_0 is the expected improvement under the current model, and EI_1 is computed using the updated model after the simulated evaluation.

Lastly, our approach does not compute the two-step acquisition value solely based on the current user's model. Instead, it leverages the combined predictions of all population models along with the current model, following the same aggregation principle as TAF. In other words, the two-step acquisition is also guided by TAF. We expand on this integration in the next section.

4 META-PO: META-LEARNING FOR PREFERENTIAL BAYESIAN OPTIMIZATION

We now propose *Meta-PO*, a novel method designed to enable efficient *preferential optimization* by leveraging prior experience from related visual design tasks. We define a visual design task as one in which a user adjusts the parameters controlling the visual appearance of a 2D photo or 3D scene to achieve a specific **target theme** (e.g., “warm,” “vintage”). Meta-PO operates on the Sequential-Gallery interaction, which is the state-of-the-art approach for interactive PBO. Using this method, users are presented with a search plane and select their preferred design candidate one

at a time. After multiple iterations of selection, the final design is determined. For the meta-learning component, Meta-PO adopts TAF, which maintains low computational complexity and operates effectively without the need for large datasets.

We consider two scenarios for meta-learning. **(1) Cross-user, same-theme:** Prior users have optimized designs for a specific theme, and the current user seeks to achieve the same theme. This setting focuses on generalizing across different users with similar goals. **(2) Cross-user, cross-theme:** Prior users have optimized for a range of themes, while the current user is pursuing a *novel, unseen* theme. This more challenging case requires Meta-PO to generalize across both users and stylistic goals.

For both scenarios, we assume a group of “prior users” has completed full optimization sessions — a phase we refer to as “population modeling.” The resulting population models are then leveraged to support the design process of “new users” during the “deployment” phase. Crucially, even users targeting the same theme may have different subjective preferences and different target images. Furthermore, themes themselves can vary widely in style, adding complexity to knowledge transfer.

Below, we first outline the key challenges of Meta-PO. We then introduce the strategies and components of Meta-PO to address these challenges, followed by a detailed explanation of its workflow. Finally, we conduct a series of simulated experiments to evaluate the effectiveness of Meta-PO in preferential optimization tasks.

4.1 Challenges for Meta-PO

While Liao et al. [36] used TAF to improve sample efficiency in human-in-the-loop optimization, their method does not address the specific challenges of preferential optimization. To develop Meta-PO for Sequential-Gallery-based interactions, we identify key challenges unique to preference-based tasks.

The first challenge is **how to intelligently adjust the weights of population models based on ongoing observations**. Prior work computed weights solely from predictive uncertainty (i.e., variance) of each GP model [36], reflecting confidence in prior data but remaining fixed during optimization. This is limiting for preferential tasks, where user preferences are highly subjective and vary across individuals and themes. Static weights can cause irrelevant models to dominate. To support effective meta-learning, weights must adapt in real-time, emphasizing knowledge from prior models that align with the new user's unique preferences.

The second challenge, specific to Sequential-Gallery-based interaction, is **how to select the third point to construct a search plane**. Constructing the 2D plane requires at least three candidates. Following the original method, we retain the current best design (i.e., the most preferred so far) as the first point. The second point is selected via meta-learning (TAF), which aggregates population models to identify the highest acquisition value. However, the third point remains an open problem. The original work [26] addresses this by constructing a line orthogonal to the first two points and searching along it. While this promotes diversity, it restricts exploration to a limited subspace, possibly missing promising regions.

The final challenge is **how to flexibly balance the influence of population models and the current model over time**. This

reflects a broader trade-off: population models provide guided exploration but may suppress personalization, while the current model supports personalization but lacks prior guidance. Though TAF internally balances these weights, the dynamics are not directly controllable or linked to practical factors like the remaining search budget. In realistic scenarios, one may want to lean on prior models early and shift toward personalization as the budget tightens.

4.2 Core Elements and Design Principles of Meta-PO

Meta-PO consists of two elements: a gallery of *population models* which are gathered from prior users, and a *current model* which is constructed when deploying on the end-users. We adopted TAF, which is a weighted-sum approach to aggregate acquisition values from population models and the current model to support and guide the ongoing optimization.

4.2.1 Core Elements. Population models: Population models are built from the optimization experiences of prior users. For each optimization for a specific user in population modeling, we train a GP model based on user evaluations. Each model captures explored regions and observed outcomes, enabling effective reuse of prior knowledge. These models are stored as a gallery and dynamically aggregated during deployment for the new users.

Current model: During deployment, for each new user, we construct a dedicated GP model using only data from the ongoing task. This model captures user-specific characteristics and evolves as new observations are gathered.

4.2.2 Design Principles. To intelligently adjust the weights of population models based on ongoing observations, we depart from prior work that assigns weights solely based on predictive uncertainty (i.e., variance) from each model [36]. Instead, we use a similarity-based weighting mechanism that dynamically adapts to the current user’s observed preferences. Specifically, we compute the similarity between the current model and each population model, assigning higher weights to those that align more closely. The computation of model similarity will be detailed in the following section. This ensures that relevant prior experience exerts stronger influences, while less relevant models are appropriately downweighted, minimizing the risk of misleading the optimization.

Another principle, specific to Sequential-Gallery interaction, is the **selection of the third design point** needed to construct a 2D search plane. Prior work addresses this via a heuristic: it defines a line orthogonal to the first two selected points and searches along this line [26]. However, this approach does not incorporate information from prior users and restricts exploration to a predefined subspace. In contrast, we propose a more principled strategy by combining TAF with a two-step acquisition function [68], which explicitly accounts for the impact of future evaluations.

The last design principle is **to enable more flexible control of the trade-off between the weights on the population models and the current model.** We employ a time-dependent decay factor that gradually reduces the influence of population models over the course of the optimization. This allows population models to guide exploration in the early stages—when data from the current

user is scarce—and enables the current model to take over as more personalized observations become available.

4.3 Method

Following our design principles, we propose our Meta-PO. In this section, we first detail how the TAF is integrated into preferential optimization, and further construct a search plane (corresponding to the first two design principles). Then, we detail how we introduce the decay factor to dynamically balance the population models and the current model (corresponding to the third principle).

4.3.1 Transfer Acquisition Function-Ranks (TAF-R). A key design principle of Meta-PO is to adaptively adjust the weights of the models based on observations. Rather than relying on static model confidence as in prior work [36], we adopt TAF-R (R stands for “Ranks” of model similarities), a variant that updates model weights adaptively based on the relative ranking of model similarities. Model similarity is measured by how well one model’s predictions align with those of another.

Specifically, TAF-R monitors all design points that the current user has compared (e.g., between x_1 and x_2). For each such pair, both the current model and each population model predict which option is preferred. If a population model’s prediction aligns with the current user model (e.g., both prefer x_1 over x_2), its rank increases; if not, it decreases. These ranks are then used to compute similarity-based weights. Finally, these weights are used to aggregate acquisition values as defined in Equation 7, ensuring that population models more consistent with the user’s evolving preferences exert greater influence during optimization. We refer readers to the original paper [67] for additional implementation details.

4.3.2 Constructing Search Plane with Two-Step Acquisition Function. To construct the 2D search space for Sequential-Gallery, we require three distinct design points. In addition to retaining the current best design as the first point (denoted as x_k^*), we must determine two additional candidates. We utilize TAF-R to construct the second point by identifying the design with the highest acquisition value under the similarity-weighted ensemble of prior models. This point is denoted as $x_{k,1}^{AF}$. While this is analogous to how Sequential-Gallery selects the second point using a single-model acquisition function, our approach leverages TAF-R to aggregate information across multiple models for more informed guidance.

To further select the third point, we integrate TAF-R with the 2-step acquisition function, as described in Equation 3.4. This point is denoted as $x_{k,2}^{AF}$. This approach simulates hypothetical evaluations of the second point, updates the GP accordingly, and computes the expected improvement for all potential third candidates based on this updated model. In the context of preferential tasks, to perform Monte Carlo sampling similar to the approach in the original paper [68]. Specifically, we use GP to predict the mean and variance of the second query point ($x_{k,1}^{AF}$) and all points that formed the search plane in the previous iteration. We then sample from their predicted distributions and compare these sampled values to generate hypothetical comparison outcomes. These hypothetical comparisons are then used to update the GP model. By doing so, it anticipates how the selection of the second point will influence future evaluations, allowing the third point to complement the search plane

with greater diversity and long-term utility. This strategy leads to more principled exploration than heuristic-based search.

The search plane is constructed using x_k^+ as the center, with $x_{k,1}^{AF}$, $x_{k,2}^{AF}$, and their symmetric counterparts (reflected across x_k^+) forming the corners. From this plane, a new set of candidate designs is sampled and presented to the user.

4.3.3 Balancing the Population Models and the Current Model via the Decay Factor. To balance the influence of prior population models and the current user model, we introduce a time-dependent decay factor that adjusts their contributions throughout the deployment. Prior models provide useful guidance early on but may not reflect the current user's preferences. As more user-specific data accumulates, the current model becomes more informative. This motivates a dynamic weighting strategy: population models guide early exploration, while the current model dominates in later stages. Following Liao et al. [36], we define a decay function $d(k)$ that adjusts weights based on the current iteration k , parameterized by d_1 and d_2 to control decay duration and slope:

$$d(k) = \begin{cases} 1, & \text{if } k \leq d_1 \\ 1 - (k - d_1)d_2, & \text{if } d_1 < k \leq d_1 + \frac{1}{d_2} \\ 0, & \text{otherwise} \end{cases} \quad (9)$$

This decay factor is applied to each population model's weight: $w_i^+ = d(k) \cdot w_i$, allowing their influence to gradually diminish as the optimization progresses. Thus, the overall acquisition function at iteration k , denoted $a_k(x)$, is then defined as:

$$a_k(x) = (1 - d(k)) \cdot \mathbb{E}[I_s(x)] + \sum_{i=1}^M w_i^+ \cdot I_i(x) \quad (10)$$

Here, $\mathbb{E}[I_s(x)]$ is the expected improvement computed from the current model, and $I_i(x)$ is the acquisition value from the i -th population model.

4.4 Workflow

The full workflow involving Meta-PO has two main phases: *population modeling* and *deployment*. In the population modeling phase, prior users each complete a full preference-based optimization session using the Sequential-Gallery interaction. Each session results in a GP model that captures the user's preference function. During deployment, the goal is to optimize for a new user by efficiently leveraging population models via Meta-PO.

4.4.1 Population Modeling. The initial search plane is constructed using three randomly selected points from the parameter space. Once the user provides the first selection from this candidate set, Meta-PO updates the GP model using this observation, as described in Equation 3.1.

In subsequent iterations, the three design points that define the search plane are determined using the acquisition function. Consider iteration k as an example. Based on all accumulated observations up to this point, denoted $\mathcal{D}_k$, the system selects the most preferred design—i.e., the design chosen by the user in the previous iteration—as the first point, x_k^+ .

The second point is determined by the acquisition function $a_k(\cdot)$. Specifically, the design with the highest acquisition value across the search space $\mathcal{X}$ is selected as:

$$x_{k,1}^{AF} = \arg \max_{x \in \mathcal{X}} a_k(x) \quad (11)$$

To determine the third point, we employ a two-step acquisition strategy. This method anticipates the effect of evaluating the second point and selects the next best candidate under this updated belief. In the context of preferential observations, we assume the second point $x_{k,1}^{AF}$ is preferred over all prior observations, i.e., $x_{k,1}^{AF} \succ \mathcal{D}_k$. Under this assumption, we augment the dataset as follows:

$$\mathcal{D}'_k = \mathcal{D}_k \cup \{x_{k,1}^{AF} \succ \mathcal{D}_k\} \quad (12)$$

Using $\mathcal{D}'_k$, we update the GP hyperparameters and re-estimate the similarity weights of the population models, thereby recomputing the acquisition function $a'_k(x)$. The third point is then selected based on:

$$x_{k,2}^{AF} = \arg \max_{x \in \mathcal{X}} a_k^{2\text{-step}}(x) \quad (13)$$

The search plane then takes x_k^+ as the center, with $x_{k,1}^{AF}$, $x_{k,2}^{AF}$, and their counterparts as corners. The user explores the options on this plane and selects the design that best aligns with their preferences. This completes one iteration of the optimization loop. The updated preference observation is then added to $\mathcal{D}_k$, and the three key points are recomputed to construct the next search plane. After all population models are gathered, they are stored as a gallery to support the deployment of Meta-PO.

4.4.2 Deployment. When deploying Meta-PO on new users, the same overall procedure is followed. However, the first search plane is initialized slightly differently: only the center point (x_k^+) is sampled randomly, while the remaining two points ($x_{1,1}^{AF}$ and $x_{1,2}^{AF}$) are determined using TAF and two-step TAF. In this initial stage, all population models are assigned equal weights, and the current model is assigned zero weight, as no user-specific data has yet been collected. The search plane is then constructed using the same procedure described in the population modeling phase. In subsequent iterations, the selection of the three design points follows the same structure as in the population modeling phase. However, the acquisition functions used in deployment differ: rather than being computed solely from the current model, the acquisition function $a_k(x)$ is now computed using TAF, which aggregates predictions from the population models based on dynamically updated weights. Similarly, the two-step acquisition function $a_k^{2\text{-step}}(x)$ is computed using the two-step TAF formulation (Equation 3.4). The rest of the interaction follows the population modeling workflow: the user explores the 2D search plane, selects the most preferred design, and the model updates based on the new observation. This iterative loop continues until the optimization concludes. The mechanism is illustrated in Figure 2.

4.5 Simulated Tests

We evaluated Meta-PO through simulated tests to assess its effectiveness and compare design choices (see Appendix A). We focused on two components: (1) strategies for weighting population




Figure 2: Meta-PO’s working mechanism during deployment. Meta-PO first stores a set of population models (shown in the gray area). During deployment, it constructs a 2D search plane per iteration using three distinct design points in a high-dimensional design space. The first point, x_k^+ , is the most preferred design so far and serves as the plane’s center. The second point, $x_{k,1}^{AF}$, is selected by aggregating acquisition values from both the population models and the current model using TAF-R. The third point, $x_{k,2}^{AF}$, is derived from a two-step extension of TAF-R. Meta-PO then reflects these two points across x_k^+ to create $x_{k,1}'^{AF}$ and $x_{k,2}'^{AF}$. Jointly, these four points define the full 2D plane; Meta-PO then samples a grid of candidates within it and presents them to the user. The user’s selection updates the current model (GP_u), which in turn allows TAF-R to adjust model weights based on the similarity between GP_u and each population model GP_i . A new plane is constructed in the next iteration using these updated weights and new acquisition values.

models and (2) methods for selecting the third point to form the search plane. We simulated user preferences by randomly shifting and scaling benchmark functions. Non-meta baselines were used to construct population models for meta-learning variants. All meta-learning methods improved both convergence and final performance compared to non-transfer baselines.

On average, meta-learning methods required only 5–10 iterations to reach the final performance achieved by non-meta baselines. At termination, they achieved 85.86% lower regret compared to non-meta methods. Among the search strategies, the two-step acquisition method consistently outperformed orthogonal exploration, reducing final regret by an average of 29.89% across all test functions. For model weighting, the similarity-based TAF-R approach outperformed the confidence-based TAF-M, achieving a 51.06% improvement in average final regret. The combination of TAF-R and two-step acquisition yielded the best overall performance and is used as the default configuration in our user studies.

5 INTERACTION

To evaluate the effectiveness of Meta-PO in improving preferential optimization efficiency for practical applications, we implemented two applications: image enhancement and virtual scene lighting design. Image enhancement is a common task in everyday life, requiring users to tweak multiple parameters to achieve the desired visual outcome. Similarly, virtual scene lighting design involves modifying lighting parameters to control the appearance of a scene in Virtual Reality. Through these applications, we aim to

demonstrate Meta-PO’s applicability in real-world scenarios and its potential benefits for end-users. We conducted user studies to evaluate Meta-PO’s performance in both **cross-user transfer** and **cross-theme generalization**.

5.1 Image Enhancement

Image enhancement is a widely used feature in social media platforms. We implemented an image enhancement application based on Sequential-Gallery with Meta-PO to enhance the optimization process. We then conducted a user study to compare the performance and user experience between Meta-PO and Sequential-Gallery.

5.1.1 Implementation. Our application includes 12 adjustable parameters, including contrast, saturation, brightness, and color adjustments for shadows, midtones, and highlights across the RGB channels, which remain the same as used in the original Sequential-Gallery. Figure 3 illustrated the user interface. In each iteration, users were presented with a gallery of candidate images and selected the one that best matched their preference. Upon proceeding to the next iteration, a new set of candidates was generated based on their previous selection. Users could terminate the optimization process and export their final design when satisfied. Notably, the instructions shown in Figure 3 (e.g., “make it cold”) were provided for the user study but were not part of the standard user workflow.

5.1.2 User Study Design. We conducted a user study using our image enhancement application to investigate whether Meta-PO can effectively transfer optimization experience across both users




Figure 3: Image enhancement interface. Users are presented with 25 candidate images and select the one that best aligns with their desired visual effect (‘theme’). Our user study interface added the theme-based instruction (top right) and buttons to choose whether they were satisfied with the current iteration’s outcome or needed to see additional candidates.

and themes. To achieve this, we adopted a between-group study design with three conditions:

- *No-Transfer*: Participants used the Sequential-Gallery interaction to optimize images without leveraging prior optimization experience. This condition not only served as the baseline, but also as the participants’ optimization data forming population models used in the other two conditions. We employed the two-step lookahead acquisition function for the Sequential-Gallery method, ensuring that the resulting models could serve as population models for subsequent conditions.
- *Cross-users, same theme (Cross-Users)*: Participants in this condition optimized images to achieve the same themes as those in the *No-Transfer* condition. However, their optimization process was accelerated using Meta-PO and population models provided by *No-Transfer* participants. By comparing results between this condition and *No-Transfer*, we evaluated the effectiveness of Meta-PO in cross-user transfer.
- *Cross-themes and cross-users (Cross-Themes)*: Participants in this condition optimized images to achieve different themes than those in the previous two conditions, also using Meta-PO and population models provided by *No-Transfer* participants. By comparing their results to *No-Transfer*, we assessed the effectiveness of Meta-PO in cross-theme generalization.

We selected image themes from existing Instagram and TikTok filters. The *No-Transfer* and *Cross-Users* conditions used: warm, cold, golden hour, and blue hour. The *Cross-Themes* condition used: winter, summer, pastel, and vintage.

We generated 100 landscape images using Stable Diffusion [50], all from the same prompt¹. In each trial, one image was randomly selected, and participants were asked to adjust it to a target theme within 15 iterations. In each iteration, they selected the most preferred candidate; if satisfied, they clicked “satisfactory,” otherwise “see more” to continue. To assess optimization efficiency, we recorded the *minimum number of iterations* needed to reach satisfaction. We also collected subjective feedback using the Creativity Support Index (CSI) [9], which measures Enjoyment, Exploration, Expressiveness, Engagement, and Satisfaction on a 7-point Likert scale (1 = strongly agree, 7 = strongly disagree).

5.1.3 Procedure. Participants first gave informed consent and completed a demographic questionnaire. They were then introduced to the study and completed a warm-up trial with a randomly selected image and a given theme. Participants were given the theme names and instructed to interpret each one freely. The themes were presented solely as abstract terms without examples or additional guidance for two reasons: first, all themes were selected from popular social media platforms, and all participants reported

¹“River landscape with clear water and surrounding forest.”

prior experience using them to share photos; second, this approach was intended to encourage open-ended interpretation. This setup reflects real-world scenarios in which users aim to achieve a general visual effect without a specific target image. Once ready, they proceeded to the main study. In each trial, a random image was shown, and participants adjusted it to match a given theme. Each trial lasted 15 iterations, regardless of satisfaction, though participants could indicate satisfaction at any point by clicking a "satisfied" button. Each participant completed 16 trials (four themes $\times$ four images), with trial order randomized. After all trials, participants filled out the CSI questionnaire. The study lasted approximately 45 minutes.

5.1.4 Participants & Apparatus. We recruited a total of 36 participants, with 12 participants in each group. In *No-Transfer*, participants consisted of 7 males and 5 females, aged 25–29 ($M = 27.17$, $SD = 1.59$). In *Cross-Users*, the participants included 8 males and 4 females, aged 24–30 ($M = 27.75$, $SD = 1.81$). Similarly, *Cross-Themes* included 7 males and 5 females, aged 25–31 ($M = 27.92$, $SD = 1.88$).

The study was conducted on a desktop running Windows 10, equipped with a 12th Gen Intel Core i7 CPU and an NVIDIA GeForce RTX 2080 GPU. Participants interacted with the user interface using a mouse and keyboard, with the interface displayed in full-screen mode on a 27-inch Philips monitor (model: 275V8LA).

We implemented both the population and current models as Gaussian Processes using the automatic relevance determination (ARD) Matérn 5/2 kernel, following the recommendation of Snoek et al. [54]. The kernel hyperparameters are determined via maximum a posteriori (MAP) estimation using L-BFGS, a gradient-based local optimization method [41]. To find the optimal acquisition values, we employed DIRECT, a global optimization algorithm [23], as multiple local optima may exist. We ran the algorithm for 50 iterations to ensure thorough exploration of the search space. We implemented the decay factor using $d_1 = 3$ and $d_2 = 8$, based on findings from a pilot study.

5.1.5 Results. We analyzed the *least iteration* required for achieving satisfactory and subjective feedback across conditions. We conducted Shapiro-Wilk tests for the *least iteration* and results suggested that the normality assumption was not violated ($p > .05$). Therefore, we performed a one-way ANOVA to analyze the effect of *Condition* on it. For post-hoc analysis, we employed independent t-tests with Bonferroni correction. For the qualitative feedback, we applied a Kruskal-Wallis test to each of the five metrics and we used Mann-Whitney U tests with Bonferroni correction for post-hoc pairwise comparisons.

Meta-PO could help participants to achieve their desired image adjustments using prior optimization experience on different participants and themes. We first analyzed whether and when participants found their satisfied image adjustments in all three conditions. Across all trials, participants in *No-Transfer* reported achieving a satisfactory result within 15 iterations in 76.57% of trials. In contrast, participants in *Cross-Users* and *Cross-Themes* reached a satisfactory outcome in 97.40% and 95.83% of trials, respectively. We then analyzed the *least iteration* count required for participants to achieve a satisfactory result, excluding trials in which participants failed to reach a satisfactory outcome. Statistical analysis revealed a significant difference in the *least iteration*

across groups ($[F_{2,33} = 18.37, p < 0.001]$). Post-hoc comparisons indicate that *Cross-Users* participants ($M = 5.86$, $SD = 1.20$) required significantly fewer iterations to reach a satisfactory result compared to *No-Transfer* participants ($M = 9.54$, $SD = 2.19$, $t(11) = 5.11$, $p < 0.001$). This finding suggests that Meta-PO successfully accelerates the optimization process using prior experience on different human participants. Similarly, *Cross-Themes* participants ($M = 7.41$, $SD = 1.28$), who adjusted images to a different theme, also required fewer iterations compared to *No-Transfer* participants ($t(11) = 3.23$, $p < 0.01$). This suggests that the optimization experience can transfer across different themes as well. However, post-hoc comparisons further indicate that *Cross-Themes* participants required more iterations than *Cross-Users* participants ($t(11) = -3.87$, $p < 0.01$), which is expected given that the target theme for *Cross-Themes* differs from the optimization experience stored in the population models.

The intermediate results shown in Figure 4 illustrate the convergence process in both *No-Transfer* and *Cross-Users*. Additionally, Figure 5 presents the intermediate results within 5 iterations for *Cross-Users* and *Cross-Themes*, demonstrating that participants were able to achieve satisfactory results within 5 iterations by leveraging prior optimization experience.

Participants perceived greater exploration potential when using Meta-PO. Statistical analysis revealed a significant main effect of group on *Exploration* ($\chi^2(2) = 10.53$, $p < 0.01$). Post-hoc comparisons indicate that participants in *Cross-Users* ($M = 2.33$, $SD = 0.49$) reported a greater ability to explore diverse ideas during the optimization process compared to *No-Transfer* participants ($M = 3.33$, $SD = 0.89$, $W = 0.83$, $p < 0.01$). Similarly, participants in *Cross-Themes* ($M = 2.58$, $SD = 0.63$) also perceived a greater exploratory capacity compared to *No-Transfer* participants ($W = 0.76$, $p < 0.05$). These findings suggest that Meta-PO enabled broader exploration during optimization, potentially due to our two-step acquisition function-based search space construction, which allows participants to explore a more diverse search space within the same number of iterations.

Participants reported higher satisfaction with their image adjustment results when using Meta-PO. Statistical analysis revealed a significant main effect of group on *Satisfaction* ($\chi^2(2) = 19.77$, $p < 0.001$). Post-hoc comparisons indicate that *Cross-Users* participants ($M = 2.00$, $SD = 0.60$) reported significantly higher satisfaction with their results compared to *No-Transfer* participants ($M = 3.92$, $SD = 0.90$, $W = 0.94$, $p < 0.001$). Similarly, *Cross-Themes* participants ($M = 2.21$, $SD = 0.54$) also reported greater satisfaction compared to *No-Transfer* participants ($W = 0.93$, $p < 0.001$). These findings suggest that Meta-PO enabled a more efficient convergence toward participants' desired adjustments, allowing them to leverage more iterations to refine their designs, ultimately leading to more satisfactory results.

5.1.6 Discussion. These results suggest that even partial transfer of prior experience can effectively reduce user effort and improve their experience in interactive design tasks. By lowering the number of interactions required to reach satisfactory outcomes, Meta-PO makes PBO more practical for casual, low-effort use cases. This usability improvement broadens the applicability of PBO to everyday design tasks—such as color selection, font customization, and




Figure 4: Comparison of the intermediate results generated by Meta-PO and by Sequential Gallery (baseline) across four themes. Participants reached satisfactory results in fewer iterations with Meta-PO than with the Sequential Gallery.




Figure 5: Intermediate results after the 1st, 3rd, and 5th iterations using Meta-PO across all eight themes. Leveraging previous optimization experience, participants achieved suitable results within five iterations.




Figure 6: CSI questionnaire results for image enhancement across all three conditions (means and standard deviations). (*: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$).

digital fabrication, where users may not have the time or expertise to engage in prolonged optimization.

5.2 Preliminary Study on Virtual Scene Lighting Design

Beyond image enhancement, we also implemented a virtual scene lighting design application, allowing participants to explore the search space continuously using the joystick of the controllers. This application demonstrates that Meta-PO also supports more flexible exploration and interaction beyond gallery-based search. Similar to the image enhancement study, we conducted a preliminary user study with a total of 10 participants (across both population modeling and deployment) to demonstrate the effectiveness of Meta-PO in this use case.

5.2.1 Design. In this study, we included a single virtual living room scene and asked participants to design the lighting to achieve different visual themes. This setup simulates a realistic application where users could customize their home lighting in VR before making physical adjustments.

Our study design closely follows that of the image enhancement study. Participants were divided into two groups: *No-Transfer* served as the population users, utilizing the Sequential-Gallery interaction with continuous search space without meta-learning, while *Meta-PO* used Meta-PO, leveraging prior optimization data from *No-Transfer*.

For *No-Transfer*, participants were asked to adjust the lighting to achieve warm and cold themes. For *Meta-PO*, in addition to warm and cold, they also designed for golden hour and pastel lighting effects, which are different from the population participants.

We included a single light source, with participants adjusting its 3D position, RGB values, and intensity, resulting in a seven-parameter search space. Each participant performed 15 iterations per theme and was asked to indicate the first iteration where they felt satisfied with the result. After completing all themes, participants filled out the Creativity Support Index (CSI) questionnaire,

evaluating the system across the same five key perspectives as in the previous study.

5.2.2 Procedure. The procedure closely follows the previous study, with one key difference: instead of selecting a preferred image from 25 candidates, participants explored the search space freely using the joystick on their controllers. Once they were satisfied with their current design, they pressed a button on the controller to proceed to the next iteration. During the study, a small square with a cursor was displayed in mid-air in front of the participant, illustrating the exploration space. This visualization helped participants better understand the boundary and identify unexplored areas.

5.2.3 Participants & Apparatus. We recruited 10 participants in total, with 4 participants assigned to *No-Transfer* and 6 participants to *Meta-PO*. We included fewer participants in the population model for two reasons. First, the number of parameters in this study was lower than in the previous study, resulting in a smaller search space. We also aimed to evaluate whether Meta-PO could still effectively facilitate subsequent optimization tasks with a reduced population model.

No-Transfer consisted of 3 males and 1 female, aged 25–29 ($M = 27.17$, $SD = 1.59$). *Meta-PO* included 4 males and 2 females, aged 24–30 ($M = 27.17$, $SD = 1.59$).

Participants used a Meta Quest 3 headset connected to a desktop computer with an Intel Core i7-12700K processor, an NVIDIA GeForce GTX 2080 GPU, and 32 GB of RAM. They interacted with the application using the Meta Quest 3 controllers. The optimizer implementation remained consistent with the previous study.

5.2.4 Results. Similar to the previous study, we analyzed both the *least iteration* required for achieving satisfactory and subjective feedback across groups. The primary goal of this preliminary study is to demonstrate the generalizability of Meta-PO. Therefore, we did not perform statistical analyses on the study results.

Across all trials in both groups, every participant successfully found a satisfactory design within 15 iterations. This may be attributed to the lower number of adjustable parameters compared to the previous study, as 15 iterations were sufficient to 7 seven parameters. In *No-Transfer*, participants achieved their desired effect within an average of 7 iterations ($M = 6.85$, $SD = 1.48$). While participants in *Meta-PO*, where prior optimization experience was leveraged, reached a satisfactory result in approximately 4 iterations ($M = 4.05$, $SD = 1.39$) when tasked with achieving the same visual effect as *No-Transfer*. When *Meta-PO* participants aimed to adjust the virtual lighting to new themes, they required an average of 5 iterations ($M = 5.33$, $SD = 1.47$). These results demonstrate that Meta-PO effectively accelerates visual appearance optimization in 3D spaces, even when relying on a smaller set of prior optimization experiences from population models.

Figure 7 illustrated the intermediate results using Meta-PO in all four themes. These results demonstrated that participants could achieve satisfactory results within 5 iterations.

In qualitative feedback, participants in *Meta-PO* ($M = 2.33$, $SD = 0.52$) reported feeling that they could explore a wider range of possibilities compared to those in *No-Transfer* ($M = 3.38$, $SD = 0.48$). Additionally, they expressed higher satisfaction with their final lighting design ($M = 2.42$, $SD = 0.49$) than participants in




Figure 7: Meta-PO-based optimization of immersive lighting appearances: Intermediate results for four themes.




Figure 8: CSI questionnaire results for virtual lighting appearance across the two conditions (means and standard deviations).

No-Transfer ($M = 3.63$, $SD = 0.48$). These findings align with our previous study, further supporting the effectiveness of Meta-PO in enabling participants to explore more diverse design possibilities while achieving greater satisfaction with their results.

6 DISCUSSION

Meta-PO accelerates the visual appearance optimization process by transferring prior users' preferences. At its core, Meta-PO follows

TAF-R to compute the similarities between the current user's specific model and all the previously gathered population models and then dynamically adjusts the weights that determine which prior model has a stronger influence to support the current optimization. As Meta-PO builds on the interaction of Sequential-Gallery, we must derive a third design point. Unlike the original work's approach to searching along the orthogonal line, Meta-PO integrates the two-step acquisition function with TAF-R to identify the third point. This simulates the potential future impact of a second design point before selecting the third, offering a broader exploration within the design space.

We evaluated Meta-PO on two visual design tasks: image enhancement and virtual lighting design. In the image enhancement task, participants achieved their desired results in fewer than 6 iterations on average, requiring significantly fewer iterations than baseline methods that do not utilize prior experience. Even when applying optimization experiences across different themes, participants reached their goals in under 8 iterations on average, outperforming the baseline. Additionally, participants reported that Meta-PO allowed them to explore more possibilities and ultimately achieve more satisfying results. In the virtual lighting design task, participants similarly found that Meta-PO enabled them to reach their desired outcomes more quickly. These findings suggest that Meta-PO can infer users' implicit preferences and selectively transfer prior optimization experiences without requiring strong assumptions or domain-specific knowledge. We believe Meta-PO has the

potential to assist designers and users in adjusting visual appearance settings more efficiently, and to support other tasks where user preferences are central.

The impact of the number of population models. Our study showed that Meta-PO accelerates visual appearance optimization by leveraging prior knowledge, specifically the preference models and optimization experiences of 12 participants across 192 images. However, we did not examine the impact of the number of population models on performance. In principle, a larger collection of population models increases the chances of finding one or more that closely resemble the current user’s preferences, thereby improving the accuracy and relevance of the guidance during early-stage optimization. Prior work supports this idea [36] and showed that additional models can boost sample efficiency by offering richer, more diverse prior knowledge—but the benefit gradually diminishes after a certain point. The optimal number is highly task-dependent.

Another key consideration is the trade-off between leveraging prior knowledge and maintaining computational efficiency. Using more population models increases guidance but also raises computational costs at deployment time, potentially impacting real-time interaction and user experience. This scalability issue is a shared limitation of TAF-based methods. Future work could explore strategies such as clustering similar models or using sparse Gaussian Processes to reduce inference overhead. More advanced directions include replacing the weighted-sum formulation with scalable Bayesian neural networks, which can leverage larger prior datasets without increasing inference cost.

Open questions for transferring experience across models and themes. Our study demonstrated that Meta-PO enables effective knowledge transfer both across user models and across visual themes. In particular, even when a user’s target theme was unseen during population modeling, transfer was still beneficial, primarily because the new theme shared stylistic similarities with existing ones. However, the effectiveness of transfer is highly dependent on the degree of similarity between themes and user preferences. This raises several open questions. Can we estimate the “distance” between a new theme and previously modeled themes before deployment, in order to anticipate the effectiveness of transfer? How can we quantify similarity in a way that’s meaningful for guiding preference-based optimization? Future work should consider addressing these questions, potentially based on semantic distances and by analyzing labeled images. Furthermore, the effectiveness of Meta-PO under entirely novel or highly abstract themes remains uncertain. In such cases, stylistic similarity may be ambiguous or absent altogether, limiting the relevance of prior models. This underscores the need for mechanisms that can detect when population knowledge is no longer applicable and adapt accordingly. In more extreme cases, meta-learning may not only be unhelpful; it may even hinder the optimization process. For example, if the population models primarily capture warm-related styles, they may offer misleading suggestions when the user is targeting a cold or high-contrast theme. Future work should explore mechanisms to detect when negative transfer is occurring and automatically fall back to optimizing from scratch, or selectively prune misleading population models based on early interaction signals.

Selecting relevant themes for initializing optimization. In our current implementation, Meta-PO leverages all available prior models across themes when operating in cross-theme scenarios. However, as discussed earlier, incorporating models from semantically irrelevant themes may result in negative transfer, even hindering optimization performance. A potential direction for future work is to intelligently select a subset of relevant themes to initialize the optimization. One potential approach is to measure semantic similarity (e.g., textual description) and visual similarity (e.g., visual features) between the current task and prior themes, and restrict initialization to those that are most relevant. This targeted selection could potentially improve convergence and enhance generalization when the model base increases.

Alternative implementation of meta-learning for PBO. While Meta-PO builds on TAF-R and a two-step acquisition function, alternative meta-learning mechanisms for Preferential Bayesian Optimization (PBO) could be explored. One direction is to adopt a kernel-based approach, where meta-learning is incorporated by designing task-specific or theme-specific kernels, similar to Brochu et al. [4]. Such kernels could allow smoother generalization across users or themes without maintaining an explicit set of population models. Another direction involves Bayesian Neural Networks (BNNs), which offer amortized inference and better scalability as the number of prior tasks increases. Additionally, one could consider meta-learned acquisition functions that adapt exploration strategies based on prior optimization behavior [61], or few-shot preference modeling that learns shared representations using neural encoders. While Meta-PO focuses on a model-based ensemble approach, future work could compare these alternatives to understand trade-offs in scalability, generalization, interpretability, and performance.

Preference Shifts During Interaction and Re-adaptation. In practice, users’ preferences may evolve during the course of optimization. As participants interact with diverse visual candidates, their goals may shift—for example, from initially seeking a “subtle warmth” to later favoring a “dramatic vintage” effect. This drift presents a challenge: early observations may no longer reflect the user’s current intent, leading to degraded performance if the optimizer continues to treat all past feedback equally. To address this, future work should consider more principled models for handling non-stationary preferences. For example, the Gaussian Process could be modified to include time-decay kernels, where older data points are downweighted as newer ones arrive — an idea explored in non-stationary BO literature. Alternatively, a more intuitive solution is to introduce a sliding window over the most recent iterations or use forgetting factors that diminish the influence of outdated selections. These strategies would allow the model to stay responsive to emerging user preferences while maintaining sample efficiency. Integrating such mechanisms into PBO could make systems like Meta-PO more robust in dynamic or exploratory design tasks.

Incorporating user ratings and rankings with preference feedback. Meta-PO generates preference models by asking users to select their favorite option from a set of candidates, implicitly generating multiple pairwise comparisons between the chosen design and the others. While effective, future work could further increase sample

efficiency by incorporating richer forms of user feedback. For example, asking users to rank their top K options would yield a richer set of pairwise comparisons. Alternatively, collecting scalar ratings or preference scores could provide fine-grained signals across all or multiple candidates. These additional forms of feedback could significantly enhance the optimizer’s ability to model user preferences and accelerate convergence. However, such extensions will also introduce increased cognitive load. Future work should carefully examine this trade-off, balancing the gains in sample efficiency against the additional effort required from users.

Generalization of Meta-PO to Other Applications. Although Meta-PO is developed for visual appearance design, its underlying method, meta-learned preference-based optimization, can generalize to a wide range of human-in-the-loop design tasks. Many domains require users to iteratively refine high-dimensional parameters solely based on subjective criteria, without an explicit or measurable objective function. Examples include sound design, haptic interface personalization, lighting adjustment, game design, and interaction technique calibration, where user preferences drive decision-making and optimization efficiency is crucial. In such scenarios, leveraging prior user data through population models and dynamically adapting to new preferences, as done in Meta-PO, can accelerate convergence and reduce user effort. Future work could explore how well our Meta-PO or broader meta-learning for PBO can generalize across modalities and application contexts.

Comparison to foundation models driven by prompt iteration. Recent advances in foundation models, such as large language and vision-language models, provide alternative means of interactive design. Users can steer visual outcomes by iteratively editing prompts in tools like Stable Diffusion or ChatGPT, gradually approaching their desired aesthetic. In this process, users must iteratively optimize and refine their prompts while navigating various trade-offs and requirements, such as providing explicit objective goals with verbal instructions, balancing exploration and exploitation by users themselves [51, 69]. By contrast, Meta-PO offers a principled, sample-efficient optimization approach that is explicitly grounded in user feedback. It makes user preferences tractable through Gaussian Process modeling, supports transparent selection logic, and enables interpretable control over trade-offs across multiple criteria. Users directly influence the optimization process through comparisons, not abstract linguistic cues. Future work may explore ways to bring these benefits into prompt-based workflows; for instance, by learning preference-aware prompt templates or integrating interactive optimization loops with foundation models.

7 CONCLUSION

We introduce Meta-PO, a novel computational method that improves the efficiency of Preferential Bayesian Optimization (PBO) by leveraging prior optimization experiences through meta-learning. It infers prior users’ implicit preferences, stores them as models, and leverages them to suggest design candidates for new users. By dynamically weighting prior preference models based on their relevance to the current user, Meta-PO facilitates more efficient preferential optimization. Moreover, we integrate the two-step acquisition function approach to construct a 2D search plane in a

high-dimensional design space. We evaluated the effectiveness of Meta-PO through a series of simulated tests and two user studies in both image enhancement and virtual reality lighting optimization. Our results demonstrate that Meta-PO significantly accelerates the optimization process both in **cross-user transfer** and **cross-theme generalization**. These results showed that Meta-PO improved the optimization efficiency on both familiar and unseen design goals. By enabling efficient, scalable, and generalizable preference-driven optimization, Meta-PO brings personalized visual optimization closer to end-users, paving the way for intuitive and adaptive design tools that could be integrated into everyday creative workflows. Finally, this work showcases how data-driven approaches can enable more generalizable and efficient human-in-the-loop Bayesian optimization. We anticipate this work to shed light on the border application scope of scaling up optimization by prior knowledge across users and tasks.

8 OPEN SCIENCE

The implementation of Meta-PO is available on our project page at https://siplab.org/projects/Visual_Appearance_Optimization.

ACKNOWLEDGMENTS

Yi-Chi Liao was supported by the ETH Zurich Postdoctoral Fellowship Programme. Zhipeng Li was partially supported by the Swiss National Science Foundation (Grant No. 10004941).

REFERENCES

- [1] Moritz Bächer, Emily Whiting, Bernd Bickel, and Olga Sorkine-Hornung. 2014. Spin-it: Optimizing moment of inertia for spinnable objects. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–10.
- [2] Tianyi Bai, Yang Li, Yu Shen, Xinyi Zhang, Wentao Zhang, and Bin Cui. 2023. Transfer learning for bayesian optimization: A survey. *arXiv preprint arXiv:2302.05927* (2023).
- [3] Rémi Bardenet, Mátyás Brendel, Balázs Kégl, and Michele Sebag. 2013. Collaborative hyperparameter tuning. In *International conference on machine learning*. PMLR, 199–207.
- [4] Eric Brochu, Tyson Brochu, and Nando De Freitas. 2010. A Bayesian interactive optimization approach to procedural animation design. In *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. 103–112.
- [5] Bilal Catkin and Volkan Patoglu. 2023. Preference-based human-in-the-loop optimization for perceived realism of haptic rendering. *IEEE Transactions on Haptics* 16, 4 (2023), 470–476.
- [6] Liwei Chan, Yi-Chi Liao, George B Mo, John J Dudley, Chun-Lien Cheng, Per Ola Kristensson, and Antti Oulasvirta. 2022. Investigating positive and negative qualities of human-in-the-loop optimization for designing interaction techniques. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [7] Jonlin Chen, Masaru Ishii, Kristin L Bater, Halley Darrach, David Liao, Pauline P Huynh, Isabel P Reh, Jason C Nellis, Anisha R Kumar, and Lisa E Ishii. 2019. Association between the use of social media and photograph editing applications, self-esteem, and cosmetic surgery acceptance. *JAMA facial plastic surgery* 21, 5 (2019), 361–367.
- [8] Yi Fei Cheng, Christoph Gebhardt, and Christian Holz. 2023. Interactionadapt: Interaction-driven workspace adaptation for situated virtual reality environments. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–14.
- [9] Erin Cherry and Celine Latulipe. 2014. Quantifying the creativity support of digital tools through the creativity support index. *ACM Transactions on Computer-Human Interaction (TOCHI)* 21, 4 (2014), 1–25.
- [10] Mark Colley, Pascal Jansen, Mugdha Keskar, and Enrico Rukzio. 2025. Improving External Communication of Automated Vehicles Using Bayesian Optimization. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 80, 16 pages. <https://doi.org/10.1145/3706598.3714187>
- [11] Nigel Cross. 2004. Expertise in design: an overview. *Design studies* 25, 5 (2004), 427–441.

- [12] Ritendra Datta, Dhiraj Joshi, Jia Li, and James Z Wang. 2006. Studying aesthetics in photographic images using a computational approach. In *Computer Vision–ECCV 2006: 9th European Conference on Computer Vision, Graz, Austria, May 7–13, 2006, Proceedings, Part III* 9. Springer, 288–301.
- [13] James R Deneault, Woojae Kim, Jiseob Kim, Yuzhe Gu, Jorge Chang, Benji Maruyama, Jay I Myung, and Mark A Pitt. 2025. Preferential Bayesian optimization improves the efficiency of printing objects with subjective qualities. *Digital Discovery* 4, 3 (2025), 723–737.
- [14] Yubin Deng, Chen Change Loy, and Xiaou Tang. 2017. Image aesthetic assessment: An experimental survey. *IEEE Signal Processing Magazine* 34, 4 (2017), 80–106.
- [15] Yan Duan, John Schulman, Xi Chen, Peter L. Bartlett, Ilya Sutskever, and Pieter Abbeel. 2016. RL\$^2\$: Fast Reinforcement Learning via Slow Reinforcement Learning. CoRR abs/1611.02779 (2016). arXiv:1611.02779 <http://arxiv.org/abs/1611.02779>
- [16] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*. PMLR, 1126–1135.
- [17] Peter I Frazier. 2018. A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811* (2018).
- [18] David Ginsbourger and Rodolphe Le Riche. 2010. Towards Gaussian process-based optimization with finite time horizon. In *mODA 9–Advances in Model-Oriented Design and Analysis: Proceedings of the 9th International Workshop in Model-Oriented Design and Analysis held in Bertinoro, Italy, June 14–18, 2010*. Springer, 89–96.
- [19] Javier González, Zhenwen Dai, Andreas Damianou, and Neil D Lawrence. 2017. Preferential bayesian optimization. In *International Conference on Machine Learning*. PMLR, 1282–1291.
- [20] Bing-Jing Hsieh, Ping-Chun Hsieh, and Xi Liu. 2021. Reinforced few-shot acquisition function learning for bayesian optimization. *Advances in Neural Information Processing Systems* 34 (2021), 7718–7731.
- [21] Kuo-Ying Huang. 2009. Challenges in human-computer interaction design for mobile devices. In *Proceedings of the World Congress on Engineering and Computer Science*, Vol. 1. San Francisco, USA, 1–6.
- [22] Koki Iwai, Yusuke Kumagai, Yuki Koyama, Masahiro Hamasaki, and Masataka Goto. 2025. Constrained Preferential Bayesian Optimization and Its Application in Banner Ad Design. *arXiv preprint arXiv:2505.10954* (2025).
- [23] Donald R Jones, Cary D Perttunen, and Bruce E Stuckman. 1993. Lipschitzian optimization without the Lipschitz constant. *Journal of optimization Theory and Applications* 79 (1993), 157–181.
- [24] Andreas Karrenbauer and Antti Oulasvirta. 2014. Improvements to keyboard optimization with integer programming. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*. 621–626.
- [25] Mohammad M. Khajah, Brett D. Roads, Robert V. Lindsey, Yun-En Liu, and Michael C. Mozer. 2016. Designing Engaging Games Using Bayesian Optimization. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 5571–5582. <https://doi.org/10.1145/2858036.2858253>
- [26] Yuki Koyama, Issei Sato, and Masataka Goto. 2020. Sequential gallery for interactive visual design optimization. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 88–1.
- [27] Yuki Koyama, Issei Sato, Daisuke Sakamoto, and Takeo Igarashi. 2017. Sequential line search for efficient visual design optimization by crowds. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–11.
- [28] Andreas Krause and Cheng Ong. 2011. Contextual gaussian process bandit optimization. *Advances in neural information processing systems* 24 (2011).
- [29] Remi Lam, Karen Willcox, and David H Wolpert. 2016. Bayesian optimization with a finite budget: An approximate dynamic programming approach. *Advances in Neural Information Processing Systems* 29 (2016).
- [30] Bryan Lawson. 2006. *How designers think*. Routledge.
- [31] Dingzeyu Li, David IW Levin, Wojciech Matusik, and Changxi Zheng. 2016. Acoustic voxels: Computational optimization of modular acoustic filters. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–12.
- [32] Wendi Li, Ting Wang, and Wing WY Ng. 2021. Population-based hyperparameter tuning with multitask collaboration. *IEEE transactions on neural networks and learning systems* 34, 9 (2021), 5719–5731.
- [33] Yang Li, Yu Shen, Huaijun Jiang, Wentao Zhang, Zhi Yang, Ce Zhang, and Bin Cui. 2022. Transbo: Hyperparameter optimization via two-phase transfer learning. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 956–966.
- [34] Zhipeng Li, Christoph Gebhardt, Yves Inglin, Nicolas Steck, Paul Streli, and Christian Holz. 2024. Situationadapt: Contextual ui optimization in mixed reality with situation awareness via llm reasoning. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–13.
- [35] Yi-Chi Liao. 2023. Human-in-the-Loop Design Optimization. (2023).
- [36] Yi-Chi Liao, Ruta Desai, Alec M Pierce, Krista E Taylor, Hrvoje Benko, Tanya R Jonker, and Aakar Gupta. 2024. A Meta-Bayesian Approach for Rapid Online Parametric Optimization for Wrist-based Interactions. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–38.
- [37] Yi-Chi Liao, John J Dudley, George B Mo, Chun-Lien Cheng, Liwei Chan, Antti Oulasvirta, and Per Ola Kristensson. 2023. Interaction design with multi-objective bayesian optimization. *IEEE Pervasive Computing* 22, 1 (2023), 29–38.
- [38] Yi-Chi Liao, Sunjun Kim, Byungjoo Lee, and Antti Oulasvirta. 2020. Button simulation and design via fdv models. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [39] Yi-Chi Liao, George B Mo, John J Dudley, Chun-Lien Cheng, Liwei Chan, Per Ola Kristensson, and Antti Oulasvirta. 2024. Practical approaches to group-level multi-objective Bayesian optimization in interaction technique design. *Collective Intelligence* 3, 1 (2024), 26339137241241313.
- [40] Yi-Chi Liao, Paul Streli, Zhipeng Li, Christoph Gebhardt, and Christian Holz. 2025. Continual Human-in-the-Loop Optimization. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–26.
- [41] Dong C Liu and Jorge Nocedal. 1989. On the limited memory BFGS method for large scale optimization. *Mathematical programming* 45, 1 (1989), 503–528.
- [42] Lydia Manikonda, Yuheng Hu, and Subbarao Kambhampati. 2014. Analyzing user activities, demographics, social network structure and user-generated content on Instagram. *arXiv preprint arXiv:1410.8099* (2014).
- [43] Luca-Maxim Meinhardt, Clara Schramm, Pascal Jansen, Mark Colley, and Enrico Rukzio. 2025. Fly Away: Evaluating the Impact of Motion Fidelity on Optimized User Interface Design via Bayesian Optimization in Automated Urban Air Mobility Simulations. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [44] George Mo, John Dudley, Liwei Chan, Yi-Chi Liao, Antti Oulasvirta, and Per Ola Kristensson. 2024. Cooperative multi-objective Bayesian design optimization. *ACM Transactions on Interactive Intelligent Systems* 14, 2 (2024), 1–28.
- [45] Jonas Mockus. 1975. On the Bayes methods for seeking the extremal point. *IFAC Proceedings Volumes* 8, 1 (1975), 428–431.
- [46] Alex Nichol and John Schulman. 2018. Reptile: a scalable metalearning algorithm. *arXiv preprint arXiv:1803.02999* 2, 3 (2018), 4.
- [47] Peter O'Donovan, Aseem Agarwala, and Aaron Hertzmann. 2015. Designscape: Design with interactive layout suggestions. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 1221–1224.
- [48] Michael A Osborne, Roman Garnett, and Stephen J Roberts. 2009. Gaussian processes for global optimization. (2009).
- [49] Romain Prévost, Emily Whiting, Sylvain Lefebvre, and Olga Sorkine-Hornung. 2013. Make it stand: balancing shapes for 3D fabrication. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–10.
- [50] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis With Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 10684–10695.
- [51] Antonio Sabatella, Andrea Ponti, Ilaria Giordani, Antonio Candelieri, and Francesco Archetti. 2024. Prompt optimization in large language models. *Mathematics* 12, 6 (2024), 929.
- [52] Nicolas Schweighofer and Kenji Doya. 2003. Meta-learning in reinforcement learning. *Neural Networks* 16, 1 (2003), 5–9.
- [53] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. 2015. Taking the human out of the loop: A review of Bayesian optimization. *Proc. IEEE* 104, 1 (2015), 148–175.
- [54] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. Practical bayesian optimization of machine learning algorithms. *Advances in neural information processing systems* 25 (2012).
- [55] Kevin Swersky, Jasper Snoek, and Ryan P Adams. 2013. Multi-task bayesian optimization. *Advances in neural information processing systems* 26 (2013).
- [56] Shion Takeno, Masahiro Nomura, and M Karasuma. 2022. Preferential Bayesian optimization with hallucination believer. In *Proc. NeurIPS Workshop Gaussian Processes, Spatiotemporal Model, Decision-Making Syst.* 1–11.
- [57] Yuki Tatsukawa, I-Chao Shen, Mustafa Doga Dogan, Anran Qi, Yuki Koyama, Ariel Shamir, and Takeo Igarashi. 2025. FontCraft: Multimodal Font Design Using Interactive Bayesian Optimization. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [58] Kashyap Todi, Daryl Weir, and Antti Oulasvirta. 2016. Sketchplore: Sketch and explore with a layout optimiser. In *Proceedings of the 2016 ACM conference on designing interactive systems*. 543–555.
- [59] Kristi Tsukida, Maya R Gupta, et al. 2011. How to analyze paired comparison data. *Department of Electrical Engineering University of Washington, Tech. Rep. UWEETR-2011-0004* 1 (2011).
- [60] Ricardo Vilalta and Youssef Drissi. 2002. A perspective view and survey of meta-learning. *Artificial intelligence review* 18 (2002), 77–95.
- [61] Michael Volpp, Lukas P Fröhlich, Kirsten Fischer, Andreas Doerr, Stefan Falkner, Frank Hutter, and Christian Daniel. 2019. Meta-learning acquisition functions for transfer learning in bayesian optimization. *arXiv preprint arXiv:1904.02642* (2019).
- [62] Jane X Wang. 2021. Meta-learning in natural and artificial intelligence. *Current Opinion in Behavioral Sciences* 38 (2021), 90–95.

- [63] Xilu Wang, Yaochu Jin, Sebastian Schmitt, and Markus Olhofer. 2023. Recent advances in Bayesian optimization. *Comput. Surveys* 55, 13s (2023), 1–36.
- [64] Karl Weiss, Taghi M Khoshgoftaar, and DingDing Wang. 2016. A survey of transfer learning. *Journal of Big data* 3 (2016), 1–40.
- [65] John H Williamson, Antti Oulasvirta, Per Ola Kristensson, and Nikola Banovic. 2022. *Bayesian Methods for Interaction and Design*. Cambridge University Press.
- [66] Martin Wistuba and Josif Grabocka. 2021. Few-shot bayesian optimization with deep kernel surrogates. *arXiv preprint arXiv:2101.07667* (2021).
- [67] Martin Wistuba, Nicolas Schilling, and Lars Schmidt-Thieme. 2018. Scalable gaussian process-based transfer surrogates for hyperparameter optimization. *Machine Learning* 107, 1 (2018), 43–78.
- [68] Jian Wu and Peter Frazier. 2019. Practical two-step lookahead Bayesian optimization. *Advances in neural information processing systems* 32 (2019).
- [69] Zhaozhuo Xu, Zirui Liu, Beidi Chen, Yuxin Tang, Jue Wang, Kaixiong Zhou, Xia Hu, and Anshumali Shrivastava. 2023. Compress, then prompt: Improving accuracy-efficiency trade-off of llm inference with transferable prompt. *arXiv preprint arXiv:2305.11186* (2023).
- [70] Kenta Yamamoto, Yuki Koyama, and Yoichi Ochiai. 2022. Photographic lighting design with photographer-in-the-loop bayesian optimization. In *Proceedings of the 35th annual ACM symposium on user interface software and technology*. 1–11.
- [71] Dani Yogatama and Gideon Mann. 2014. Efficient transfer learning method for automatic hyperparameter tuning. In *Artificial intelligence and statistics*. PMLR, 1077–1085.

A SIMULATED TESTS

This paper introduces Meta-PO, which leverages meta-learning to utilize optimization experience from similar past tasks, thereby accelerating the optimization process for the current task. Importantly, our Meta-PO is built upon the interaction of Sequential-Gallery, which constructs a search plane at every iteration — the system has to suggest three design points to form such a plane at each iteration.

In this appendix, we conduct simulated evaluations for two primary goals. First, we aim to demonstrate the general effectiveness of enhancing PBO with meta-learning. Second, two decision factors can lead to different Meta-PO implementations, including (1) different ways of computing the weights on the prior models and (2) different ways of constructing the search plane, specifically regarding the decision of the third design point. We also aim to investigate which is the most effective combination. In detail, we have the following goals:

- (1) **Validating the effectiveness of meta-learning for preference optimization:** We investigate whether meta-learning improves preference optimization by comparing our approach to a similar method that does not utilize prior optimization experience.
- (2) **Comparing model weight assignment strategies:** We examine two distinct strategies for assigning weights to population models [67].
- (3) **Comparing different approaches for constructing the search plane:** We specifically compare using two strategies of deriving the third decision point at every iteration, which ultimately leads to different search planes.

As our primary goal is to compare the performance of different approaches in a meta-learning setting, we have to generate a group of similar yet different test functions. Each specific test function can be seen as a unique user’s implicit preference model. We simulated different user preferences by randomly shifting and scaling a given base function. All experiments were performed on a Windows 10 system equipped with a 12th Gen Intel(R) Core i7 CPU and an NVIDIA GeForce RTX 2080 GPU.

The results of these synthetic tests support our decision to employ TAF-R as the main approach to assigning model weights and using the 2-step acquisition function to determine the third design point. For more details about TAF-R and the 2-step acquisition function, please refer to subsection 4.3.

A.1 Compared Methods

Here we detail different combinations of implementing Meta-PO, which will then be compared throughout all synthetic tests.

Following the framework of TAF [67], there exist two different approaches to assigning weights to both population models and the current model. The original work suggested two approaches, which we aim to compare.

TAF-M: weight based on confidence TAF-M adjusts each model’s weight based on its confidence in the current sampling point. This confidence is estimated using the variance of the model’s predictions. A higher variance indicates greater uncertainty, suggesting that the model’s predictions are less reliable and

should be assigned a lower weight. Conversely, a lower variance implies greater confidence in the predictions, leading to a higher weight.

TAF-R: weight based on prediction ranking alignment TAF-R

adjusts the weights based on the similarity between previous models and the current model. To quantify this similarity, both the previous models and the current model generate predictions for all sampling points and rank them accordingly. A previous model receives a higher weight if its ranking of predictions closely aligns with that of the current model, indicating a greater degree of similarity. In contrast, if the rankings diverge significantly, the previous model’s influence is reduced to prevent misleading guidance in the optimization process.

We also compared two different plane decision strategies: two-step acquisition function [68] and orthogonal exploration, which is proposed in the original paper of Sequential-Gallery [26]. Note that, as introduced in subsection 4.3, we select a third point to define the search plane in addition to the current best observation, x_k^+ , and the point with the highest acquisition value, $x_{k,1}^{AF}$. In the below variations, the selection of the first two points (x_k^+ and $x_{k,1}^{AF}$) remains the same. The only difference is the selection of the third point, which is denoted as $x_{k,2}^{AF}$.

Two-step Acquisition Function This approach basically follows Wu and Frazier [68]. It estimates the value at the design point that currently has the highest acquisition value by Monte Carlo sampling. Then, it updates the GP accordingly, recomputes the acquisition values, and then selects the next point with the highest updated acquisition value as $x_{k,2}^{AF}$. For more details, please refer to section 3.

Orthogonal exploration Previous work [26] selected the third point by first computing the vector u formed by the current best observation, x_k^+ , and the point with the highest acquisition value, $x_{k,1}^{AF}$. Then, the third point, $x_{k,2}^{AF}$, is determined such that the vector v , formed by x_k^+ and $x_{k,2}^{AF}$, satisfies the orthogonality condition $u \cdot v = 0$.

Therefore, there are two ways of deriving the weights on the models and two ways of selecting the third point to form a search plane — making four possible implementations of Meta-PO in total.

Finally, we include three baselines to compare with: *No-Transfer* refers to preferential optimization without leveraging any prior knowledge. Yet, there are two approaches to constructing the search plane (two-step v.s. orthogonal exploration), leading to two variations. Additionally, we include a *Random* method, where the exploration plane is constructed based on the best observation (x^+) but explored in a random direction (the plane corners are randomly selected).

Finally, we summarize all seven methods in Table 1. These include four meta-learning variants (two weight strategies $\times$ two plane strategies), two non-transfer baselines (No-Transfer-O, No-Transfer-T), and one random baseline. For the meta-learning methods, the same plane construction strategy is used in both the population modeling and deployment phases to ensure consistency. For instance, Meta-PO-M-T, which uses the two-step acquisition strategy during deployment, builds its population models using the

Table 1: The implementations of PBO that we compare in the simulated tests.

Method Name	Meta-learning	Weight assignment strategy	Plane decision strategy
Random	No	\	\
No-Transfer-O	No	\	Orthogonal exploration
Meta-PO-M-O	Yes	TAF-M	Orthogonal exploration
Meta-PO-R-O	Yes	TAF-R	Orthogonal exploration
No-Transfer-T	No	\	Two-step acquisition function
Meta-PO-M-T	Yes	TAF-M	Two-step acquisition function
Meta-PO-R-T	Yes	TAF-R	Two-step acquisition function

same two-step method (i.e., No-Transfer-T). This alignment ensures a fair comparison and meaningful knowledge transfer across tasks.

In all simulated tests, we apply a weight decay factor with parameters $d_1 = 5$ and $d_2 = 0.1$. In each iteration, we initialize the search for the highest acquisition values with 80 randomly sampled points and perform 100 iterations for each to ensure a reasonable solution.

A.2 Benchmark Functions (Base Functions)

To evaluate the generalizability of our method, we conducted experiments using four benchmark functions: the 3D Hartmann function², the 6D Hartmann function³, Isotropic Gaussian, and Rosenbrock. The first two functions, Hartmann 3D and 6D, have multiple local optima, which results in a complex landscape that poses challenges for optimizers in locating the global optimum. The latter two illustrated that our method can be generalized to higher dimensions; we used 15-dimensional Isotropic Gaussian and 20-dimensional Rosenbrock, consistent with the setup in previous work [26].

During the optimization process, we normalize the parameter space to $[0, 1]$ and the output space to $[-1, 1]$, using regret as the optimization metric to quantify how closely the current solution approximates the global optimum. To simulate user preferential selection, in each iteration, the design point with the highest value from the search plane will be selected as the preferred design.

A.3 Generating Synthetic Population Models

In the simulated tests, to generate the synthetic users that contribute to the population models, we first implement standard PBO without leveraging any prior optimization experience. To introduce variance across synthetic users, we shift and scale the base functions to create a diverse set of synthetic user functions and their corresponding models. This is a common approach found in evaluating meta-Bayesian optimization [36, 54].

To shift the test function, we apply a shift factor to each dimension of the input parameter, defined as $x'_n = x_n + \delta_n$, where x_n is the original input parameter in each dimension, x'_n is the shifted input vector, and δ_n is the shift factor, which follows a uniform distribution:

$$\delta_n \sim U(-\text{shift_range}, \text{shift_range}). \quad (14)$$

Similarly, to scale the test function, we apply a scaling factor to the function values by directly multiplying a scalar with the objective value. The scalar s is also sampled from a uniform distribution:

$$s \sim U(1 - \text{scale_range}, 1 + \text{scale_range}). \quad (15)$$

In the following tests, we introduce variations by setting $\text{shift_range} = 0.05$ and $\text{scale_range} = 0.1$.

A.4 Test 1: 3D Hartmann

We generate 20 user functions, with each optimization method running for 30 iterations per user function. 10 of them are seen as prior users for population modeling, and the rest 10 are seen as the final target users during deployment.

A.4.1 Results. Goal 1: Validating the effectiveness of meta-learning. To fairly compare the performance with and without meta-learning in preferential optimization, we first examine the best performance achieved by each. Specifically, *Meta-PO-R-T* is compared against *No-Transfer-T* along with *Random* serving as the baseline, presented in Figure 9a. Overall, the optimal meta-learning-based condition (*Meta-PO-R-T*) achieved the overall best performance. While the No-Transfer preferential optimization using the two-step acquisition function (*No-Transfer-T*) outperforms random exploration, *Meta-PO-R-T* achieves a regret below 0.01 in fewer than five iterations, indicating nearly converging, whereas *Random* exploration requires 30 iterations to reach the same result. Furthermore, *Meta-PO-R-T* constantly yields better performance than *No-Transfer-T* throughout. Overall, meta-learning enables faster convergence and overall better performance, showcasing its effectiveness.

Goal 2: Comparing model weight assignment strategies. We compared the two different weight assignment strategies alongside the two evaluation set selection methods (Figure 9c). The results indicate that meta-learning methods based on TAF-R adapt more effectively to the current test function than those based on TAF-M. This advantage may arise from the fact that TAF-M relies on the variance of predictions at each point. However, when test functions are shifted and scaled, the population model's predictions may become misleading, even if they exhibit high confidence. In contrast, TAF-R-based methods leverage ranking alignment between the current model and population models, making them more robust to shifts and scaling in test functions.

²<https://www.sfu.ca/ssurjano/hart3.html>

³<https://www.sfu.ca/ssurjano/hart6.html>




Figure 9: Ablation study results on 3D Hartmann. The left plot compares the optimization performance of *Random*, *No-Transfer-T*, and *Meta-PO-R-T*, showing that leveraging prior optimization experience significantly accelerates convergence and improves outcomes. The middle plot contrasts *Random*, *No-Transfer-T*, and *No-Transfer-O*, which differ in search space construction strategy, highlighting that the two-step acquisition function consistently enhances optimization efficiency. The right plot compares *Random* with four meta-learning-based methods, demonstrating their effectiveness in accelerating optimization, with TAF-R-based methods performing slightly better than TAF-M-based methods.

Goal 3: Comparing different approaches for constructing the search plane. We then compared the results of the two-step acquisition function with orthogonal exploration from prior work. In Figure 9b, the results showed that while both methods constantly enable more effective exploration than *Random*, the two-step acquisition function outperforms orthogonal exploration. Similar patterns can be found in Figure 9c where all meta-learning-based methods are compared. We can see that the ones involved two-step acquisition function (*Meta-PO-R-T* and *Meta-PO-M-T*) outperform the rest in the earlier iterations, showing more efficient convergence. This advantage arises from the fact that the two-step acquisition function is a more principled solution for identifying a promising candidate over the full design space rather than limiting the search space to an orthogonal line.

A.5 Test 2: 6D Hartmann

Similar to the previous test, we generated 20 user functions, with each optimization method running for 30 iterations. Among these 20 functions, 10 were seen as prior users for population modeling, and the rest 10 were seen as the final test users.

Goal 1: Validating the effectiveness of meta-learning. Similar to the previous test, we compare the best condition with meta-learning (*Meta-PO-R-T*) and without (*No-Transfer-T*). We observe that *Meta-PO-R-T* consistently achieves lower regret than *No-Transfer-T* from the very first iteration (Figure 10a) and maintains superior performance throughout. As the dimension of the test function increases, *Meta-PO-R-T* does not immediately reach a highly optimal solution but still achieves a regret below 0.1 in 10 iterations, while *No-Transfer-T* achieves such a performance after 25 iterations. These results suggest that our meta-learning-based methods can accelerate the optimization process on even more complex test functions.

Goal 2: Comparing the weight assignment strategies. The results in Figure 10c suggest that TAF-R-based meta-learning methods outperform TAF-M-based methods, which is consistent with the findings from the previous simulated test.

Goal 3: Comparing different approaches for constructing the search plane. In Figure 10b, we observe that the two evaluation set selection methods achieve similar regret across all iterations. This could be attributed to the complexity of the 6D Hartmann test function, which has many local optima. However, examining Figure 10b, we can observe that the two-step acquisition function (*Meta-PO-R-T*) constantly results in better performance than using orthogonal line search (*Meta-PO-R-O*), highlighting the benefit of using the two-step method in the meta-learning settings.

A.6 Test 3: 15D Isotropic Gaussian

For the 15-dimensional Isotropic Gaussian test function, we generate 20 user functions, with each optimization method running for 50 iterations per user function. The increased number of iterations reflects the increased dimension of this function, requiring more iterations to converge. Similar to prior tests, 10 user functions are used for population modeling, and 10 are used as final test functions.

Goal 1: Validating the effectiveness of meta-learning. Figure 11a shows that *Meta-PO-R-T* achieves a regret below 0.1 within the first few iterations, whereas *No-Transfer-T* requires 40 iterations to reach the same level. Again, the meta-learning-based approaches constantly outperform the ones without. The *Random* condition struggles to optimize the high-dimensional parameter space, though the test function is relatively naive. This highlights the importance of acquisition functions and exploration strategies in high-dimensional optimization.

Goal 2: Comparing the weight assignment strategies. All methods with meta-learning can start from a reasonably good search point, regardless of the weight assignment strategy or evaluation selection method (Figure 11c). Moreover, the results suggest




Figure 10: Ablation study results on 6D Hartmann. The main conclusions remain consistent with the previous results. However, in the right plot, *Meta-PO-R-T* outperforms other meta-learning-based methods, highlighting its effectiveness in handling high-dimensional optimization problems.




Figure 11: Ablation study results on 15D Isotropic Gaussian. The main conclusions remain consistent with the previous results.

that meta-learning methods based on TAF-R enable better optimization performance compared to TAF-M, which aligns with previous test results.

Goal 3: Comparing different approaches for constructing the search plane. In Figure 11b, we observe that the *No-Transfer-T* method achieves better optimization performance than *No-Transfer-O* across all iterations after the initial exploration phase. Furthermore, both *No-Transfer-T* and *No-Transfer-O* exhibit similar performance to the *Random* method during the first 10 iterations, which is consistent with findings from prior work [26]. This is because optimization without meta-learning must first explore the high-dimensional parameter space before identifying promising regions for further exploration.

A.7 Test 4: 20D Rosenbrock

Similar to the previous test, we run 50 iterations on 20 user functions for each optimization method. 10 functions are used for population modeling and the rest 10 are used for final evaluation.

Goal 1: Evaluating the effectiveness of meta-learning. The Rosenbrock function has a shape resembling a valley or a banana (in 2D cases). Its relatively simple structure makes it more manageable for the optimization process. As a result, as shown in Figure 12a, *Meta-PO-R-T* achieves a regret below 0.1 within the first few iterations by leveraging *No-Transfer-T*'s experience over 50 iterations.

Goal 2: Comparing evaluation set selection methods. Similar to previous test results, Figure 12b suggests that *No-Transfer-T* performs comparably in most iterations and achieves a better final result than *No-Transfer-O* after 50 iterations. These findings indicate that our two-step acquisition function method outperforms orthogonal exploration on both complex test functions and high-dimensional functions.

Goal 3: Evaluating weight assignment strategies for population models. Figure 12c suggests that *Meta-PO-R-T* outperforms all other meta-learning methods, further demonstrating the effectiveness of combining the TAF-R weight assignment strategy with




Figure 12: Ablation study results on 20D Rosenbrock. The main conclusions remain consistent with the previous results. The high standard deviation observed in the results of meta-learning-based methods may be due to the log form of regret is very sensitive to small value changes, particularly when regret values are close to zero.

the two-step acquisition function method. All meta-learning methods exhibit a relatively large standard deviation in the logarithm of regret after 20 iterations. This is because the log transformation of regret is highly sensitive to small changes in regret values, especially when the regrets are close to zero. As a result, in some user functions, meta-learning methods can approximate the global optimum more closely, while in others, they perform less effectively, leading to a high standard deviation.

A.8 Conclusion

Based on the results from all four simulated tests, we conclude that meta-learning methods can accelerate the optimization process. Particularly, TAF-R brings constantly better performance compared to TAF-M, owing to its adaptive mechanism of determining weights based on the similarities between the current model and prior population models. Finally, the two-step acquisition function method generally results in better performance than orthogonal exploration, especially when combined with the meta-learning-based approaches. With these insights, we determine to employ TAF-R as the primary meta-learning mechanism with a 2-step acquisition function to conduct our user study.