Constrained Preferential Bayesian Optimization and Its Application in Banner Ad Design

Koki Iwai[1] , **Yusuke Kumagae**[1] , **Yuki Koyama**[2] , **Masahiro Hamasaki**[2] and **Masataka Goto**[2]
[1]Hakuhodo DY Holdings Inc.
[2]National Institute of Advanced Industrial Science and Technology (AIST)
{koki.iwai, yusuke.kumagae}@hakuhodo.co.jp, {koyama.y, masahiro.hamasaki, m.goto}@aist.go.jp

Abstract

Preferential Bayesian optimization (PBO) is a variant of Bayesian optimization that observes relative preferences (e.g., pairwise comparisons) instead of direct objective values, making it especially suitable for human-in-the-loop scenarios. However, real-world optimization tasks often involve inequality constraints, which existing PBO methods have not yet addressed. To fill this gap, we propose *constrained preferential Bayesian optimization* (CPBO), an extension of PBO that incorporates inequality constraints for the first time. Specifically, we present a novel acquisition function for this purpose. Our technical evaluation shows that our CPBO method successfully identifies optimal solutions by focusing on exploring feasible regions. As a practical application, we also present a designer-in-the-loop system for banner ad design using CPBO, where the objective is the designer's subjective preference, and the constraint ensures a target predicted click-through rate. We conducted a user study with professional ad designers, demonstrating the potential benefits of our approach in guiding creative design under real-world constraints.

1 Introduction

Preferential Bayesian optimization (PBO) [Brochu *et al.*, 2007; Gonzalez *et al.*, 2017; Koyama *et al.*, 2022] is a variant of Bayesian optimization (BO) [Shahriari *et al.*, 2016] that observes relative preferences (e.g., pairwise comparisons) instead of direct objective values. Since relative evaluations are generally considered easier, faster, and more accurate than absolute evaluations for human subjective assessments [Brochu *et al.*, 2010; Yoshida *et al.*, 2024], PBO is particularly well-suited for problems in which human preference serves as the objective function to be maximized. It has been effectively employed to implement human-in-the-loop optimization systems for visual design [Brochu *et al.*, 2007; Brochu *et al.*, 2010; Koyama *et al.*, 2017; Koyama *et al.*, 2020; Yamamoto *et al.*, 2022].

However, existing PBO methods have not been adapted to handle constrained optimization problems. Many real-world optimization tasks (e.g., product and architectural design, drug discovery, and recommender systems) involve maximizing human preferences under additional constraints. For example, in product and architectural design, users may express preferences over usability, while the design must also satisfy constraints such as durability or thermal performance, often evaluated through physical simulations. These scenarios require human-in-the-loop optimization under costly or implicit constraints, highlighting a technical gap in extending PBO to such applications.

To address this gap, we propose a new method called *constrained preferential Bayesian optimization* (CPBO), which incorporates inequality constraints into PBO. As the core of CPBO, we introduce a novel acquisition function, *expected utility of the best option with constraints* (EUBOC), which extends an existing PBO acquisition function [Lin *et al.*, 2022]. We evaluate the proposed method through simulation experiments, highlighting its ability to find optimal solutions while effectively accounting for constraints.

As a practical application of CPBO, we propose a novel *designer-in-the-loop* framework for banner ad design, where the predicted *click-through rate* (CTR) serves as a constraint. Banner ads are used to promote products or services online, and CTR, the fraction of clicks relative to impressions, is a key metric in the advertising industry [McMahan *et al.*, 2013; Chen *et al.*, 2016; Richardson *et al.*, 2007; Zhou *et al.*, 2018]. However, focusing solely on maximizing CTR can lead to designs that, while effective in capturing clicks, may be visually unappealing or even annoying, negatively affecting brand perception [Zeng *et al.*, 2020; Zeng *et al.*, 2021]. Consequently, effective banner ad design requires maximizing visual appeal while maintaining a reasonably high CTR. This task is more complex than typical visual design tasks that focus solely on subjective preferences, such as those addressed in previous work [Brochu *et al.*, 2010; Koyama *et al.*, 2020; Yamamoto *et al.*, 2022]. To evaluate its effectiveness, we conducted a user study with professional ad designers, revealing that they positively received the concept of our framework and appreciated its potential to reduce design workload.

An extended version of this paper with an appendix is available at http://arxiv.org/abs/2505.10954

Our contributions are summarized as follows.

- We propose CPBO, a novel extension of PBO that handles inequality-constrained optimization problems.
- To implement CPBO, we propose EUBOC, a new acquisition function for CPBO. We evaluate its performance through simulation experiments.
- As a practical CPBO application, we propose a novel designer-in-the-loop framework for banner ad design that integrates CTR considerations. A user study with professional ad designers validates its real-world viability.

2 Preliminaries

Before describing our proposed CPBO technique, we describe its foundation: PBO and constrained Bayesian optimization (CBO). Especially, we detail their acquisition functions: *expected utility of the best option* (EUBO) [Lin *et al.*, 2022] for PBO, and *expected improvement with constraints* (EIC) [Gardner and Kusner, 2014] for CBO.

Throughout this paper, we assume that the input variables consist of N continuous parameters, each normalized to the range of $[0, 1]$ (without loss of generality). We denote these parameters by $x_i \in [0, 1]$ $(i = 1, \ldots, N)$, and define the vector $\boldsymbol{x} = [x_1, \ldots, x_N] \in \mathcal{X}$, where $\mathcal{X} = [0, 1]^N$. We further denote the objective function by $f : \mathcal{X} \to \mathbb{R}$ and the constraint function by $c : \mathcal{X} \to \mathbb{R}$.

2.1 Preferential Bayesian Optimization

PBO is a variant of BO that observes relative preferences instead of direct objective function values. As with standard BO, the goal of PBO is to identify the global optimum of a black-box function:

$$\boldsymbol{x}^* = \arg\max_{\boldsymbol{x}\in\mathcal{X}} f(\boldsymbol{x}). \tag{1}$$

However, PBO achieves this by iteratively collecting preference data instead of observing the function value $f(\boldsymbol{x})$ directly. In this work, we assume that each observation yields a forced choice between two candidates, denoted as $d = \boldsymbol{x}^{(i)} \succ \boldsymbol{x}^{(j)}$ (i.e., the candidate $\boldsymbol{x}^{(i)}$ is preferred over the candidate $\boldsymbol{x}^{(j)}$).

PBO is often used in human-in-the-loop settings [Koyama *et al.*, 2022], where the black-box objective function represents a human preference, and the goal is to find the most preferred option through iterative human evaluations.

In PBO, a probabilistic model serves as a surrogate for the objective function f, and this model is continuously updated based on the observed preference data through Bayesian inference. Gaussian processes (GPs) [Rasmussen and Williams, 2005] are often used as surrogates owing to their flexibility; we adopt GPs for this purpose as well.

The likelihood of a preference data $d = \boldsymbol{x}^{(i)} \succ \boldsymbol{x}^{(j)}$ is often modeled using the Thurstone-Mosteller model [Chu and Ghahramani, 2005]:

$$P(d\,|\,f) = \Phi\left(\frac{f(\boldsymbol{x}^{(i)}) - f(\boldsymbol{x}^{(j)})}{\sqrt{2}\sigma}\right), \tag{2}$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution, and σ is the standard deviation of the Gaussian noise. The likelihood of multiple preference data, $\mathcal{D} = \{d_1, d_2, \ldots\}$ is given by $P(\mathcal{D}|f) = \Pi_i P(d_i|f)$. Using this probability, the surrogate model can be updated; refer to [Chu and Ghahramani, 2005] for details.

At each iteration of PBO, the next candidate point to be compared is selected by maximizing an acquisition function derived from the updated surrogate model. The acquisition function measures a candidate's effectiveness, and its design is critical to the overall performance of PBO.

EUBO [Lin *et al.*, 2022] is one of the state-of-the-art acquisition functions for PBO. It can be expressed in closed form for two search points $\boldsymbol{x}^{(i)}$ and $\boldsymbol{x}^{(j)}$ as

$$\begin{aligned}\mathrm{EUBO}(\boldsymbol{x}^{(i)},\, \boldsymbol{x}^{(j)}) &= \mathbb{E}[\max\{f(\boldsymbol{x}^{(i)}),\, f(\boldsymbol{x}^{(j)})\}] \\ &= \Delta\Phi\left(\frac{\Delta}{\sigma}\right) + \sigma\phi\left(\frac{\Delta}{\sigma}\right) + \mu^f(\boldsymbol{x}^{(j)}),\end{aligned} \tag{3}$$

where $\phi(\cdot)$ is the probability density function of the standard normal distribution, $\Delta = \mathbb{E}[f(\boldsymbol{x}^{(i)}) - f(\boldsymbol{x}^{(j)})]$, $\sigma^2 = \mathrm{Var}[f(\boldsymbol{x}^{(i)}) - f(\boldsymbol{x}^{(j)})]$, and $\mu^f(\boldsymbol{x}^{(j)}) = \mathbb{E}[f(\boldsymbol{x}^{(j)})]$. We adopt EUBO as the basis of our proposed technique because of its efficiency and computational simplicity.

2.2 Constrained Bayesian Optimization

CBO is a variant of BO designed for optimization problems with an additional inequality constraint defined as

$$c(\boldsymbol{x}) \geq \lambda, \tag{4}$$

where c is a black-box constraint function and λ denotes a threshold for the constraint. In addition to modeling the objective function, CBO typically constructs a GP surrogate for the constraint function; we adopt this approach in our work.

EIC [Gardner and Kusner, 2014] is an acquisition function for CBO. The key idea is to multiply the *expected improvement* (EI), a popular acquisition function for BO, by the probability of satisfying the constraint; that is,

$$\mathrm{EIC}(\boldsymbol{x}) = P(c(\boldsymbol{x}) \geq \lambda)\,\mathrm{EI}(\boldsymbol{x}). \tag{5}$$

We adopt EIC as the constraint-handling method in our proposed technique because of its simplicity and compatibility with EUBO. Although many constraint-handling methods exist [Hernández-Lobato *et al.*, 2015; Amini *et al.*, 2025], most are *incompatible* with EUBO, such as those requiring direct objective values or information-theoretic assumptions. Given these considerations, we find EIC to be a suitable choice.

3 Constrained Preferential Bayesian Optimization

3.1 Problem Formulation

We aim to solve the following constrained optimization problem:

$$\boldsymbol{x}^* = \arg\max_{\boldsymbol{x}\in\mathcal{X}} f(\boldsymbol{x}) \quad \text{s.t.} \quad c(\boldsymbol{x}) \geq \lambda, \tag{6}$$

where λ denotes a threshold for the constraint. Unlike standard BO, the objective function value $f(\boldsymbol{x})$ is not directly observed; instead, we observe relative preferences among multiple candidates to infer the objective function, as in PBO methods

Algorithm 1 CPBO with pairwise comparison

1: $\mathcal{H} \leftarrow \varnothing$
2: **for** $n = 1, 2, \ldots$ **do**
3: select $\boldsymbol{x}_n^{(i)}, \boldsymbol{x}_n^{(j)}$ by optimizing EUBOC:
$$\boldsymbol{x}_n^{(i)}, \boldsymbol{x}_n^{(j)} = \operatorname*{arg\,max}_{\boldsymbol{x}^{(i)},\, \boldsymbol{x}^{(j)} \in \mathcal{X}} \mathrm{EUBOC}(\boldsymbol{x}^{(i)}, \boldsymbol{x}^{(j)})$$
4: Ask the evaluator to compare $\boldsymbol{x}_n^{(i)}$ and $\boldsymbol{x}_n^{(j)}$
5: Evaluate the constraint function to obtain $c_n^{(i)}$ and $c_n^{(j)}$
6: $\mathcal{H} \leftarrow \mathcal{H} \cup \{(\boldsymbol{x}_n^{(\text{selected})} \succ \boldsymbol{x}_n^{(\text{not selected})}), (c_n^{(i)}, c_n^{(j)})\}$
7: Update surrogate models f and c based on $\mathcal{H}$
8: **end for**

(subsection 2.1). In contrast, we assume that the constraint function value $c(\boldsymbol{x})$ can be directly observed for any given $\boldsymbol{x}$ without requiring human feedback. Following the original EIC assumptions [Gardner and Kusner, 2014], we treat $c(\boldsymbol{x})$ as a black-box function that is expensive to evaluate, noise-free, and does not prevent evaluation of the objective when violated.

In practical scenarios, the objective function f often represents human preferences for a design (i.e., visual appeal), while the constraint function c captures design feasibility (e.g., performance requirements). Accordingly, we estimate f by iteratively asking human evaluators for relative preferences among multiple provided candidates, thereby capturing subjective preferences that would otherwise be difficult to quantify.

3.2 Acquisition Function: EUBOC

We propose a new acquisition function for CPBO, *expected utility of the best option with constraints* (EUBOC), which integrates EUBO with the idea of EIC. Specifically, we propose multiplying EUBO by the probability that the two points are both feasible:

$$\begin{aligned}&\mathrm{EUBOC}(\boldsymbol{x}^{(i)}, \boldsymbol{x}^{(j)})\\&= P(c(\boldsymbol{x}^{(i)}) \geq \lambda, c(\boldsymbol{x}^{(j)}) \geq \lambda)\, \mathrm{EUBO}(\boldsymbol{x}^{(i)}, \boldsymbol{x}^{(j)}).\end{aligned} \tag{7}$$

Assuming a GP for the constraint function c, the pair $(c(\boldsymbol{x}^{(i)}), c(\boldsymbol{x}^{(j)}))$ follows a bivariate normal distribution. Therefore, the probability of satisfying the inequality constraints is

$$\begin{aligned}&P(c(\boldsymbol{x}^{(i)}) \geq \lambda, c(\boldsymbol{x}^{(j)}) \geq \lambda)\\&= \int_\lambda^\infty \int_\lambda^\infty P(c(\boldsymbol{x}^{(i)}), c(\boldsymbol{x}^{(j)}))dc(\boldsymbol{x}^{(i)})dc(\boldsymbol{x}^{(j)}).\end{aligned} \tag{8}$$

Given the GP assumption, a correlation generally exists between $c(\boldsymbol{x}^{(i)})$ and $c(\boldsymbol{x}^{(j)})$. However, deriving the cumulative distribution function values analytically for a bivariate normal distribution is challenging. Therefore, we apply the following approximation, assuming that $c(\boldsymbol{x}^{(i)})$ and $c(\boldsymbol{x}^{(j)})$ are uncorrelated:

$$\begin{aligned}&\mathit{Equation}\ 8\\&\approx \int_\lambda^\infty P(c(\boldsymbol{x}^{(i)}))dc(\boldsymbol{x}^{(i)}) \int_\lambda^\infty P(c(\boldsymbol{x}^{(j)}))dc(\boldsymbol{x}^{(j)})\\&= \left(1 - \Phi\left(\frac{\lambda - \mu^c(\boldsymbol{x}^{(i)})}{\sigma^c(\boldsymbol{x}^{(i)})}\right)\right)\left(1 - \Phi\left(\frac{\lambda - \mu^c(\boldsymbol{x}^{(j)})}{\sigma^c(\boldsymbol{x}^{(j)})}\right)\right),\end{aligned} \tag{9}$$

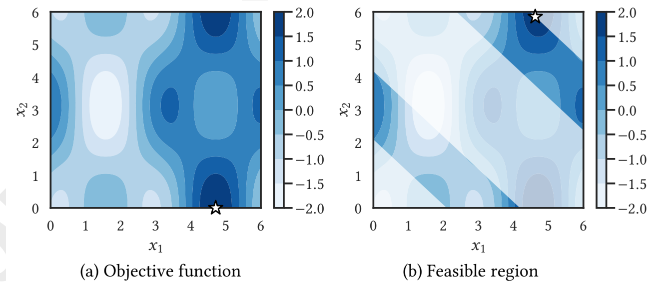

Figure 1: Two-dimensional test function used for evaluation. (a) Objective function. The star represents the maximum. (b) Constraint function overlaid onto the objective function. The lightly shaded white areas indicate infeasible regions, and the star represents the optimal solution that satisfies the constraint.

where $\mu^c(\boldsymbol{x}) = \mathbb{E}[c(\boldsymbol{x})]$ and $\sigma^c(\boldsymbol{x}) = \sqrt{\mathrm{Var}[c(\boldsymbol{x})]}$.

3.3 Algorithm

The proposed algorithm is presented in Algorithm 1, where $\mathcal{H}$ denotes the history of the evaluation data. At step n, two points are obtained by solving the maximization problem of the EUBOC acquisition function (line 3). For the obtained two points $\boldsymbol{x}_n^{(i)}$ and $\boldsymbol{x}_n^{(j)}$, we obtain the preference data $(\boldsymbol{x}_n^{(\text{selected})} \succ \boldsymbol{x}_n^{(\text{not selected})})$ from the (human) evaluator (line 4) and also their constraint values $(c_n^{(i)}, c_n^{(j)})$ by evaluating the constraint function (line 5). Then, we add these data to the history $\mathcal{H}$ and update the GP surrogate models f and c (lines 6 and 7). This process is repeated for a certain number of iterations.

3.4 Warm-Starting Constraint Surrogate Model Update (Optional)

Optionally, it is possible to pre-train the constraint surrogate model before starting the optimization iterations involving the human evaluator, given that the evaluator does not intervene in evaluating the constraint function. This *warm-start* approach facilitates starting the optimization with lower uncertainty regarding the constraints, potentially leading to improved optimization efficiency. The efficacy of this approach is demonstrated in section 4, where we pre-trained the constraint surrogate model using points randomly sampled from the search space and their constraint function values.

4 Technical Evaluation

We conducted *simulation* experiments to evaluate the proposed CPBO technique. The goal was to confirm that the proposed technique can find solutions that satisfy the inequality constraint and to understand its efficiency.

4.1 Test Function Setting

In this setting, we simulated both human responses (pairwise comparison) and constraint queries (direct observation of the constraint function values) using known synthetic test functions.

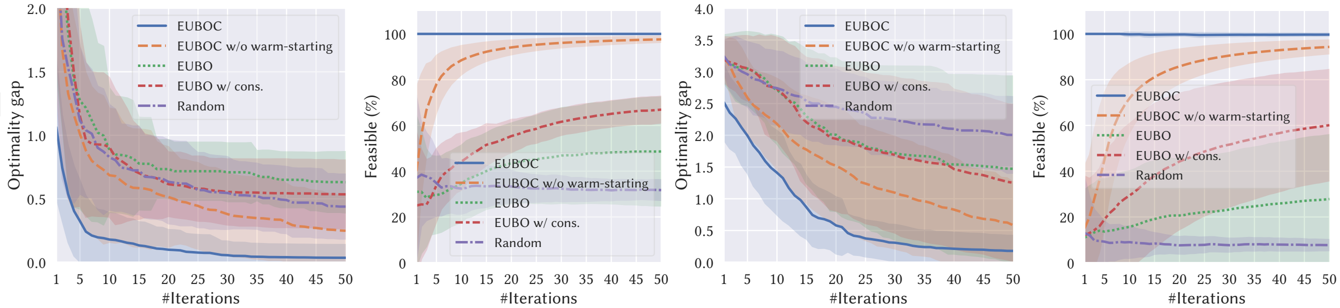

(a) 2D Optimality gap (b) 2D Feasible (c) 6D Optimality gap (d) 6D Feasible

Figure 2: Result of the test function setting. The *Optimality gap* for (a) 2D and (c) 6D test functions, where the horizontal axis represents the iteration steps and the vertical axis represents the *Optimality gap* (lower is better). The *Feasible* for (b) 2D and (d) 6D test functions, where the vertical axis represents the *Feasible* (higher is better), the lines denote the mean, and the lightly shaded areas denote the standard deviation.

Test Functions

We employed two synthetic problems for evaluation. The first comes from Gardner *et al.* [2014]. It consists of 2-dimensional (2D) objective and constraint functions composed of sine and cosine functions. Figure 1 visualizes the objective function and feasible region. The second comes from Letham *et al.* [2019]. It comprises 6-dimensional (6D) objective and constraint functions, representing a higher-dimensional case closer to real-world tasks. Its objective function is based on a Hartmann 6 function [Picheny *et al.*, 2013], while the constraint function uses the norm of $\boldsymbol{x}$ (see Appx. B.1 for details).

Methods to be Compared

We compared the following methods:

- **EUBOC**: Use our EUBOC acquisition function with warm-starting. We pre-trained the constraint surrogate model using 200 points randomly sampled from the search space.
- **EUBOC w/o warm-starting**: Use our EUBOC acquisition function without warm-starting.
- **EUBO**: Use the EUBO acquisition function. This method ignores the constraints.
- **EUBO w/ cons.**: Use the EUBO acquisition function; however, if only one of the two candidate points satisfies the constraint, we automatically treat it as the winner. This method represents a naive, post-hoc way of handling constraints.
- **Random**: Use uniform random sampling. This method ignores the constraints.

Performance Metrics

The performance metrics for the evaluation are as follows:

- *Optimality gap*: The difference between the optimal function value and the best-found function value [Wang *et al.*, 2016; Koyama *et al.*, 2020]. We set the best-found function value to the worst (smallest) objective function value if only infeasible values are observed (following [Hernández-Lobato *et al.*, 2015; Lam and Willcox, 2017]).
- *Feasible*: The cumulative proportion of sample points that satisfy the constraints up to each iteration (following [Gardner and Kusner, 2014]).

The number of iterations was set to 50. We ran each method 50 times with random initializations and recorded the mean and standard deviation of the results of all runs.

Results

Figure 2 shows the performance of each method. Overall, we find that our proposed EUBOC and EUBOC w/o warm-starting methods consistently outperformed the other methods across all settings. In addition, compared to EUBOC w/o warm-starting, EUBOC demonstrated a significant performance improvement thanks to the pre-training of the constraint surrogate model (see Appx. B.2 for further analysis).

As shown in Figure 2 (a), for the 2D test function, the EUBOC methods could quickly reduce the *Optimality gap*, whereas both EUBO and EUBO w/ cons. reduced the gap only slowly. EUBO struggled because it focused much on sampling constraint-violating regions. EUBO w/ cons., while slightly improved by considering the constraint, did not achieve significant gains, probably due to its naive constraint handling. Furthermore, EUBOC shows a smaller standard deviation, indicating robust convergence regardless of initialization.

As shown in Figure 2 (b), EUBOC successfully explored only within the feasible region throughout the iterations in the 2D setting, thanks to warm-starting and the constraint-aware acquisition function. EUBOC w/o warm-starting and EUBO w/ cons. gradually increased the *Feasible* during iterations. EUBOC w/o warm-starting, which explicitly learns the constraint, could explore the feasible region more quickly.

As shown in Figure 2 (c) and (d), for the 6D test function, EUBOC effectively focused on feasible regions in almost every iteration, efficiently reducing the *Optimality gap* compared to the other methods, similar to the performance in 2D. EUBOC w/o warm-starting also demonstrated its ability to learn the constraint during iterations and reduced the gap effectively.

4.2 Banner Ad Design Application Setting

We next simulated human responses to evaluate the effectiveness of our proposed technique in a practical banner ad design

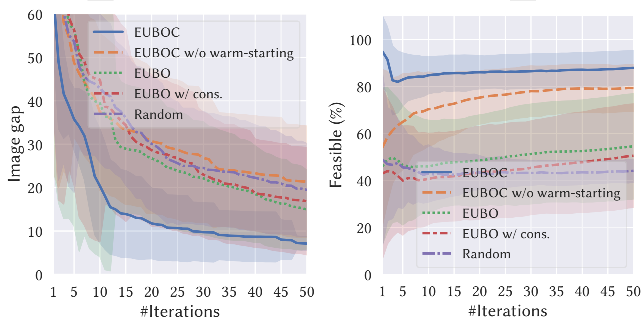

(a) Image gap

(b) Feasible

Figure 3: Result of the banner ad design application setting. (a) *Image gap* (lower is better) and (b) *Feasible* (higher is better).

context. Our application system, described in section 5.2, adjusts the colors (or layouts) of target banner ad images while ensuring a minimum required click-through rate (CTR) predicted by a machine learning model.

Task

The task was to adjust the color of a banner ad image so that it closely matched a predetermined reference image as possible. The reference image was created by recoloring an original banner ad, using parameter values selected from the feasible design space (see Appx. B.3 for the reference image). We set the threshold λ (i.e., the minimum acceptable predicted CTR) to the average of CTR values obtained from 1,000 randomly sampled parameters. The target image was described by six parameters (i.e., a 6D parameter space). At each iteration, the "human response" to any pairwise comparison was synthesized so that the chosen image would be closer to the reference image than its alternative.

Methods to be Compared

The compared methods were the same as those used in subsection 4.1, except that for the EUBOC setting, we pre-trained the constraint surrogate model using 1,000 random recoloring parameters for the warm start.

Performance Metrics

We used the *Image gap* as a performance metric [Yamamoto *et al.*, 2022]. It is defined as the average of the element-wise absolute difference between the two images, each of which is represented as a tensor in $[0, 255]^{W \times H \times 3}$ (W and H denote the width and height of the image, respectively). We regarded the *Image gap* value as 255 if only infeasible values were observed. We also recorded the *Feasible*, the proportion of the images satisfying the constraints (as described in section 4.1).

Results

Figure 3 shows the results. Our EUBOC efficiently reduced the *Image gap* by focusing its sampling on feasible regions (about 90% of the time). Consistent with subsection 4.1, EUBOC outperforms others on this problem, which resembles real-world banner ad design tasks. However, the EUBOC w/o warm-starting did not reduce the *Image gap* effectively, showing similar performance to other baselines. This is likely due to the complex shape of the constraint, which can be difficult to learn at the early stages of optimization. Warm-starting therefore proves especially beneficial here (see Appx. B.2 for further analysis). Although the constraint was not satisfied at every iteration, possibly due to the non-exhaustive pre-training of the constraint surrogate model, the *Feasible* gradually increased as the constraint surrogate model was also updated during the iterations. See Appx. B.3 for examples showing how banner ad images evolved throughout the optimization process.

5 Application: A Designer-in-the-Loop Banner Ad Design Framework

In addition to proposing CPBO as a novel optimization method, we apply it to a real-world banner ad design challenge. This demonstration not only highlights CPBO's practical value but also contributes to the human-in-the-loop design optimization field, where computational methods integrate with human subjective judgment [Brochu *et al.*, 2007; Koyama *et al.*, 2022].

5.1 Motivation and Background

Human-in-the-Loop Optimization for Graphic Design

Human-in-the-loop optimization allows human evaluators—in our case, designers—to act as the objective function and guide the search process [Koyama *et al.*, 2017; Yamamoto *et al.*, 2022; Brochu *et al.*, 2010]. Researchers have explored such methods for various design tasks [Koyama and Goto, 2022], including visual design [Koyama *et al.*, 2017; Yamamoto *et al.*, 2022; Brochu *et al.*, 2010; Koyama *et al.*, 2020; Koyama *et al.*, 2022] and interaction design [Khajah *et al.*, 2016; Liao *et al.*, 2023; Dudley *et al.*, 2019; Kadner *et al.*, 2021]. Our work adds to this body of Human-Computer Interaction (HCI) research by incorporating an additional *performance-oriented* design constraint—in our case, CTR—into the designer-in-the-loop framework.

Banner Ad Design Challenges

Banner ad design poses a unique challenge: achieving visual appeal (i.e., the designer's visual preference) while also maintaining a sufficiently high CTR. Traditional approaches may require extensive real-world testing to measure actual CTRs, which is both costly and time-consuming. Moreover, designers cannot reliably predict how visual changes affect CTR, and preference often does not align with actual CTR. Consequently, a method that *automatically ensures CTR* while letting designers focus on aesthetics is highly desirable.

Preliminary Study: Designer Preference vs. CTR

To inform and motivate our approach, we conducted a *preliminary study* with professional ad designers. Notably, the study found *no positive correlation* between a designer's preference and the actual CTR. This finding reinforces our core hypothesis that an automated system (rather than the designer) should manage CTR constraints, freeing the designer to focus on their creative intent. Detailed procedures and results of this study are provided in Appx. D.

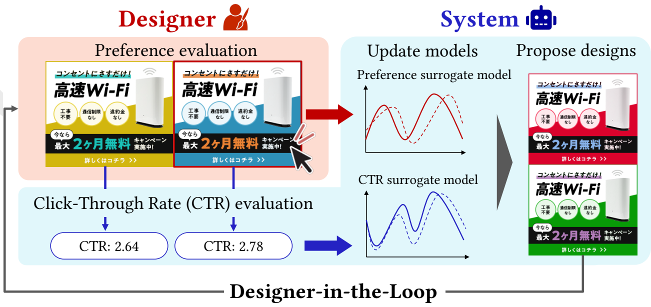

Figure 4: Concept of our designer-in-the-loop banner ad design framework. The designer provides feedback to the system about preferences, and the system predicts the CTRs. The system then updates the surrogate models and proposes new design candidates using our CPBO technique.

5.2 Design Framework

Concept

We propose a designer-in-the-loop framework for banner ad design that integrates both aesthetic preferences and CTR considerations (Figure 4). By adjusting design parameters (e.g., colors, layouts), the system aims to produce visually appealing ads while maintaining a minimum required CTR. Designers choose their preferred option from a pair of system-generated candidate ad images, focusing on creative decisions, while the system automatically handles CTR constraints using a CTR prediction model[2]. This setup addresses the challenge of simultaneously considering both preference and CTR during the design process.

This designer-in-the-loop optimization approach leverages our CPBO technique. The system maintains surrogate models for both the preference (objective) and CTR (constraint) functions. After each step of designer feedback and CTR prediction, these models are updated, and a new design candidate pair is selected to maximize preference under the CTR constraint. Through repeating this iterative process, the framework enables effective human-AI collaboration between the designer and the optimization module, ultimately producing designs that meet both aesthetic and performance goals.

System Implementation

We implemented a proof-of-concept system that supports banner ad design with two separate modes: a *color* editing mode and a *layout* editing mode. Figure 5 shows the user interface of our system in color editing mode. Our system is based on pairwise comparisons and presents the designer with two banner ad designs in each iteration. Following the designer's selection, the next two candidate designs are presented based on their preference. This iterative process allows the designer to explore more preferable designs, with the system ensuring CTR.

We implemented our CPBO technique on BoTorch [Balandat *et al.*, 2020], a BO library. We used Gaussian process models in BoTorch as the surrogate models for the objective and constraint functions. We implemented EUBOC by extending the EUBO implementation in BoTorch.

[2]Measuring actual CTRs is impractical due to the time and financial costs, so our framework uses a machine-learning-based CTR prediction model trained on real-world data instead.

Figure 5: User interface of our banner ad design system in color editing mode. This system provides pairs of design candidates for each iteration.

The CTR prediction model[3] was built using XGBoost [Chen *et al.*, 2016], trained on our dataset of real-world banner ad deployment data. This model predicts a CTR value for each given banner ad image, and those predictions update the constraint surrogate model during the optimization process (and optional warm-starting). Note that CTR prediction takes approximately 0.3 milliseconds per image.

Refer to Appx. A and C for more details.

5.3 User Study

We conducted a user study using our system (section 5.2). This experiment was carried out with the approval of the ethic examination of Research Institute of Human Engineering for Quality Life. The goal was to evaluate how our CPBO-enabled framework impacts on the user experience in banner ad design. Specifically, we aimed to assess the benefits of having a designer-in-the-loop design system responsible for CTR constraint, allowing designers to focus on their preferences. To this end, we compared two scenarios: (1) Ours: the system ensures CTR while the designer focuses on preferences, and (2) Baseline: the system does not ensure CTR.

Study Design

We recruited 11 professional ad designers (P1–P11) from an ad agency in Japan (10 females, 1 male, age: $M = 28.3, SD = 4.59$). They had an average of 5.82 years ($SD = 4.79$) of general design experience and 3.55 years ($SD = 3.01$) of specific experience in banner ad design.

We prepared two realistic banner ad images (12D and 6D parameter spaces, respectively) by asking a non-participant professional designer, using Japanese-language content. Each participant performed color-editing tasks on both images under the two system conditions. The image and system condition pairings and their conduct order were randomized. In each task, the participant selected the more preferred option from the pair of color variations the system presented, repeating this 50 times. (Note: the average time to provide preference feedback was 4.2 seconds, ranging from 1.0 to 23.1 seconds.) Participants were told to consider only their preference with

[3]Note that the CTR prediction model and the CTR surrogate model are different; the prediction model is a fixed model, while the surrogate model is dynamically updated during iterations and used for CPBO computation.

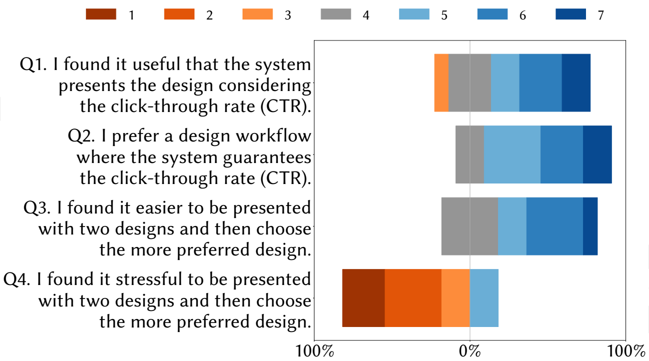

Figure 6: Questionnaire results, showing the distribution of answers for each questionnaire item. Q1 to Q3 are better to the right, and Q4 is better to the left.

our system, and both preference and CTR with the baseline. To reduce the bias from the trust in the model, we informed participants that the CTR prediction model was validated.

After finishing both tasks, participants completed a 7-point Likert-scale questionnaire (7: "strongly agree") and could provide free-form comments. We also conducted semi-structured interviews[4] regarding the tasks performed and questionnaire answers.

Questionnaire Results

Figure 6 shows the questionnaire results. Overall, professional ad designers evaluated our framework positively. Q1–Q3 focus on the CTR consideration (higher is better), while Q4 measures perceived stress (lower is better). Affirmative responses (i.e., 5–7 for Q1–Q3; 1–3 for Q4) were 63.6% (Q1), 81.8% (Q2), 63.6% (Q3), and 81.8% (Q4). Notably, Q1 and Q2, which address how the system manages CTR, received strong support.

Interview Results

We summarize the feedback on the following two points.

How the system's CTR consideration is helpful We received various positive comments about the system's CTR consideration. P5 appreciated how it accounts for CTR during the design process because *"as a designer, I want to know what makes an effective banner from a third-party [objective] perspective."* P3 noted that the system's help *"reduced the effort needed to consider CTR,"* suggesting an overall reduction in mental workload. Others also commented on differences in candidate quality compared to the baseline. P8 thought that the candidates provided by our system *"had better visibility"* and *"avoided eyestrain"*. P10 complained that, when trying the baseline, *"I got ridiculously bad ones [design candidates] many times,"* making choices trivial; in contrast, with our system, *"I was quite indecisive"* since both candidates were often good. These comments suggest that the CTR consideration helped to provide more reasonable and meaningful design candidates.

[4] Interviews were conducted in the participants' native language (Japanese); the quoted remarks here are translated.

How the system's CTR consideration changes future ad design P9 appreciated the use of actual measured data for training, saying *"It is really powerful for us [ad designers] to have the designs [created with ours] backed up by thousands of data."* P9 added that being able to *"explain this [mechanism of CTR consideration] [to the client]"* is powerful, making the tool *"branded"* and become *"a persuasive material when pitching to clients."* These remarks highlight not only the practical benefits for designers, but also the potential to improve client communication and trust, thus adding business value.

Lessons Learned

In our user study, we evaluated the benefits of our approach, where the system takes responsibility for CTR considerations in the design process, assisting ad designers. Feedback from the questionnaires and interviews showed that professional ad designers responded positively to the framework. Participants noted that it could be particularly helpful to those struggling to account for CTR or seeking an objective perspective on their designs.

Participants also indicated that incorporating CTR considerations improved the quality of design candidates in the iterative process compared to when CTR was not considered. This suggests that including the CTR constraint enhances optimization performance, supporting our technical evaluation of CPBO performance (section 4).

Finally, participants highlighted that our approach helps explain designs to clients, as they are based on actual data rather than potentially unreliable intuition. This strengthens communication and trust between designers and clients, demonstrating the business value of data-driven insights in improving both the design process and client relationships.

6 Discussions and Future Work

6.1 Improving Search Efficiency and Capability

Our experiments used search spaces of up to 12 dimensions and 50 iteration steps—enough to observe optimization behavior. In practice, designers may wish to adjust more design elements, leading to even higher-dimensional search spaces, and also minimize the required iterations. Thus, improving search efficiency is crucial.

High-dimensional BO is known to be challenging, and various methods have been proposed [Wang *et al.*, 2016; Binois and Wycoff, 2022; Long *et al.*, 2024; Hvarfner *et al.*, 2024]; future work should explore ways to combine these methods with CPBO. Another promising avenue is to enable designers to compare more than two search points simultaneously in each step [Koyama *et al.*, 2017; Koyama *et al.*, 2020; Nguyen *et al.*, 2021]. Although our EUBOC currently only supports the evaluation of two search points at a time, extending it by incorporating the concept of qEUBO [Astudillo *et al.*, 2023], an EUBO extension capable of sampling multiple search points simultaneously, would be a valuable research direction.

Another future direction is to handle categorical variables. While our EUBOC focuses on continuous inputs, we believe it can *theoretically* be extended to accommodate categorical variables—for example, by incorporating kernel adaptations for GPs [Garrido-Merchán and Hernández-Lobato, 2020].

6.2 Toward Practical Design Systems

The purpose of developing our application system was to investigate the potential benefits of a CPBO-driven design framework. The next step is to build a more comprehensive design system for professional use. Feedback from professional designers during our interviews provided multiple suggestions for improving the system. They expressed a desire to adjust not only colors or layouts, but also other design elements such as text width, font size, and letter kerning. Some designers wanted to compare more than two design candidates at a time to make a better decision (as discussed in subsection 6.1). In addition, some mentioned that, in their experience, the color and layout directions are often predetermined to some extent (e.g., by client requests) before design exploration begins, suggesting the need for a feature that limits the search space to accommodate these prior intentions, thereby enhancing the design process.

7 Conclusion

This paper proposed CPBO and a new acquisition function, EUBOC, to enable this. Our technical evaluation showed that our method efficiently reduces the gap toward optimal solutions by focusing on feasible regions. As a practical CPBO application, we proposed a designer-in-the-loop framework for designing banner ads that integrates CTR considerations. The user study demonstrated that our framework effectively reduced the design burden and proved its usefulness as a real-world CPBO application.

Ethical Statement

There are no significant ethical concerns regarding our CPBO. However, the use of CTR models in banner ad design poses potential risks, such as generating overly attention-grabbing designs that may unnecessarily encourage user clicks. Our framework mitigates these risks by facilitating collaboration between designers and the algorithm, allowing for designs that align with human subjective preferences. While this reduces the likelihood of harmful outcomes, further efforts to establish safeguards and ethical guidelines would enhance the robustness of such systems in the future.

Acknowledgements

We would like to thank Kazuhito Nishimura and his team from Hakuhodo DY ONE Inc. for their generous participation in our user study and for providing valuable insights that greatly contributed to this work.

References

[Amini *et al.*, 2025] Sasan Amini, Inneke Vannieuwenhuyse, and Alejandro Morales-Hernández. Constrained Bayesian optimization: A review. *IEEE Access*, 13:1581–1593, 2025.

[Astudillo *et al.*, 2023] Raul Astudillo, Zhiyuan Jerry Lin, Eytan Bakshy, and Peter I Frazier. qEUBO: A decision-theoretic acquisition function for preferential Bayesian optimization. In *Proc. AISTATS*, 2023.

[Balandat *et al.*, 2020] Maximilian Balandat, Brian Karrer, Daniel R Jiang, Samuel Daulton, Benjamin Letham, Andrew Gordon Wilson, and Eytan Bakshy. BoTorch: A framework for efficient monte-carlo Bayesian optimization. In *Proc. NeurIPS*, 2020.

[Binois and Wycoff, 2022] Mickaël Binois and Nathan Wycoff. A survey on high-dimensional Gaussian process modeling with application to Bayesian optimization. *ACM Trans. Evol. Learn. Optim.*, 2(2):1–26, 2022.

[Brochu *et al.*, 2007] Eric Brochu, Nando De Freitas, and Abhijeet Ghosh. Active preference learning with discrete choice data. In *Proc. NeurIPS*, 2007.

[Brochu *et al.*, 2010] Eric Brochu, Tyson Brochu, and Nando De Freitas. A Bayesian interactive optimization approach to procedural animation design. In *Proc. SCA*, 2010.

[Chen *et al.*, 2016] Junxuan Chen, Baigui Sun, Hao Li, Hongtao Lu, and Xian-Sheng Hua. Deep CTR prediction in display advertising. In *Proc. MM*, 2016.

[Chu and Ghahramani, 2005] Wei Chu and Zoubin Ghahramani. Preference learning with Gaussian processes. In *Proc. ICML*, 2005.

[Dudley *et al.*, 2019] John J Dudley, Jason T Jacques, and Per Ola Kristensson. Crowdsourcing interface feature design with Bayesian optimization. In *Proc. CHI*, 2019.

[Gardner and Kusner, 2014] Jacob R Gardner and Matt J Kusner. Bayesian optimization with inequality constraints. In *Proc. ICML*, 2014.

[Garrido-Merchán and Hernández-Lobato, 2020] Eduardo C Garrido-Merchán and Daniel Hernández-Lobato. Dealing with categorical and integer-valued variables in Bayesian optimization with Gaussian processes. *Neurocomputing*, 380:20–35, 2020.

[Gonzalez *et al.*, 2017] Javier Gonzalez, Zhenwen Dai, Andreas Damianou, and Neil D Lawrence. Preferential Bayesian optimization. In *Proc. ICML*, 2017.

[Hernández-Lobato *et al.*, 2015] José Miguel Hernández-Lobato, Michael A Gelbart, Matthew W Hoffman, Ryan P Adams, and Zoubin Ghahramani. Predictive entropy search for Bayesian optimization with unknown constraints. In *Proc. ICML*, 2015.

[Hvarfner *et al.*, 2024] Carl Hvarfner, Erik Orm Hellsten, and Luigi Nardi. Vanilla Bayesian optimization performs great in high dimensions. In *Proc. ICML*, 2024.

[Kadner *et al.*, 2021] Florian Kadner, Yannik Keller, and Constantin Rothkopf. AdaptiFont: Increasing individuals' reading speed with a generative font model and Bayesian optimization. In *Proc. CHI*, 2021.

[Khajah *et al.*, 2016] Mohammad M Khajah, Brett D Roads, Robert V Lindsey, Yun-En Liu, and Michael C Mozer. Designing engaging games using Bayesian optimization. In *Proc. CHI*, 2016.

[Koyama and Goto, 2022] Yuki Koyama and Masataka Goto. BO as Assistant: Using Bayesian optimization for asynchronously generating design suggestions. In *Proc. UIST*, 2022.

[Koyama *et al.*, 2017] Yuki Koyama, Issei Sato, Daisuke Sakamoto, and Takeo Igarashi. Sequential line search for efficient visual design optimization by crowds. *ACM Trans. Graph.*, 36(4):1–11, 2017.

[Koyama *et al.*, 2020] Yuki Koyama, Issei Sato, and Masataka Goto. Sequential gallery for interactive visual design optimization. *ACM Trans. Graph.*, 39(4):88:1–88:12, 2020.

[Koyama *et al.*, 2022] Yuki Koyama, Toby Chong, and Takeo Igarashi. Preferential Bayesian optimisation for visual design. In John H Williamson, Antti Oulasvirta, Per Ola Kristensson, and Nikola Banovic, editors, *Bayesian Methods for Interaction and Design*, chapter 8, pages 239–258. Cambridge University Press, 2022.

[Lam and Willcox, 2017] Remi Lam and Karen Willcox. Lookahead Bayesian optimization with inequality constraints. In *Proc. NeurIPS*, 2017.

[Letham *et al.*, 2019] Benjamin Letham, Brian Karrer, Guilherme Ottoni, and Eytan Bakshy. Constrained Bayesian optimization with noisy experiments. *Bayesian Anal.*, 14(2):495–519, 2019.

[Liao *et al.*, 2023] Yi-Chi Liao, John J Dudley, George B Mo, Chun-Lien Cheng, Liwei Chan, Antti Oulasvirta, and Per Ola Kristensson. Interaction design with multi-objective Bayesian optimization. *IEEE Pervasive Comput.*, 22(1):29–38, 2023.

[Lin *et al.*, 2022] Zhiyuan Jerry Lin, Raul Astudillo, Peter I Frazier, and Eytan Bakshy. Preference exploration for efficient Bayesian optimization with multiple outcomes. In *Proc. AISTATS*, 2022.

[Long *et al.*, 2024] Luo Long, Coralia Cartis, and Paz Fink Shustin. Dimensionality reduction techniques for global Bayesian optimisation. In *Proc. OPT*, 2024.

[McMahan *et al.*, 2013] H Brendan McMahan, Gary Holt, D Sculley, Michael Young, Dietmar Ebner, Julian Grady, Lan Nie, Todd Phillips, Eugene Davydov, Daniel Golovin, Sharat Chikkerur, Dan Liu, Martin Wattenberg, Arnar Mar Hrafnkelsson, Tom Boulos, and Jeremy Kubica. Ad click prediction: a view from the trenches. In *Proc. KDD*, 2013.

[Nguyen *et al.*, 2021] Quoc Phong Nguyen, Sebastian Tay, Bryan Kian Hsiang Low, and Patrick Jaillet. Top-k ranking Bayesian optimization. In *Proc. AAAI*, 2021.

[Picheny *et al.*, 2013] Victor Picheny, Tobias Wagner, and David Ginsbourger. A benchmark of kriging-based infill criteria for noisy optimization. *Struct. Multidiscip. Optim.*, 48(3):607–626, 2013.

[Rasmussen and Williams, 2005] Carl Edward Rasmussen and Christopher K I Williams. *Gaussian processes for machine learning*. Adaptive Computation and Machine Learning series. MIT Press, 2005.

[Richardson *et al.*, 2007] Matthew Richardson, Ewa Dominowska, and Robert Ragno. Predicting clicks: estimating the click-through rate for new ads. In *Proc. WWW*, 2007.

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

[Wang *et al.*, 2016] Ziyu Wang, Frank Hutter, Masrour Zoghi, David Matheson, and Nando De Feitas. Bayesian optimization in a billion dimensions via random embeddings. *J. Artif. Intell. Res.*, 55:361–387, 2016.

[Yamamoto *et al.*, 2022] Kenta Yamamoto, Yuki Koyama, and Yoichi Ochiai. Photographic lighting design with Photographer-in-the-Loop Bayesian optimization. In *Proc. UIST*, 2022.

[Yoshida *et al.*, 2024] Shigeo Yoshida, Yuki Koyama, and Yoshitaka Ushiku. Toward AI-mediated avatar-based telecommunication: Investigating visual impression of switching between user- and AI-controlled avatars in video chat. *IEEE Access*, 12:113372–113383, 2024.

[Zeng *et al.*, 2020] Eric Zeng, Tadayoshi Kohno, and Franziska Roesner. Bad news: Clickbait and deceptive ads on news and misinformation websites. In *Proc. ConPro*, 2020.

[Zeng *et al.*, 2021] Eric Zeng, Tadayoshi Kohno, and Franziska Roesner. What makes a “bad” ad? user perceptions of problematic online advertising. In *Proc. CHI*, 2021.

[Zhou *et al.*, 2018] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. Deep interest network for click-through rate prediction. In *Proc. KDD*, 2018.