

# Black-Box Optimization with Implicit Constraints for Public Policy

Wenqian Xing<sup>1</sup>, JungHo Lee<sup>2</sup>, Chong Liu<sup>3</sup>, Shixiang Zhu<sup>2</sup>

<sup>1</sup>Stanford University

<sup>2</sup>Carnegie Mellon University

<sup>3</sup>University at Albany, State University of New York

wxing@stanford.com, junghol@andrew.cmu.edu, cliu24@albany.edu, shixianz@andrew.cmu.edu

## Abstract

Black-box optimization (BBO) has become increasingly relevant for tackling complex decision-making problems, especially in public policy domains such as police redistricting. However, its broader application in public policymaking is hindered by the complexity of defining feasible regions and the high-dimensionality of decisions. This paper introduces a novel BBO framework, termed as the Conditional And Generative Black-box Optimization (CageBO). This approach leverages a conditional variational autoencoder to learn the distribution of feasible decisions, enabling a two-way mapping between the original decision space and a simplified, constraint-free latent space. The CageBO efficiently handles the implicit constraints often found in public policy applications, allowing for optimization in the latent space while evaluating objectives in the original space. We validate our method through a case study on large-scale police redistricting problems in Atlanta, Georgia. Our results reveal that our CageBO offers notable improvements in performance and efficiency compared to the baselines.

## Introduction

In recent years, black-box optimization (BBO) has emerged as a critical approach in addressing complex decision-making challenges across various domains, particularly when dealing with objective functions that are difficult to analyze or explicitly define. Unlike traditional optimization methods that require gradient information or explicit mathematical formulations, black-box optimization treats the objective function as a “black box” that can be queried for function values but offers no additional information about its structure (Pardalos et al. 2021). This methodology is especially valuable for designing public policy, such as police redistricting (Zhu et al. 2020; Zhu, Wang, and Xie 2022), site selection for emergency service systems (Xing and Hua 2022), hazard assessment (Xie et al. 2021) and public healthcare policymaking (Chandak et al. 2020). Policymakers and researchers in these areas frequently encounter optimization problems embedded within complex human systems, where decision evaluations are inherently implicit, and conducting them can be resource-intensive. Black-box optimization, therefore, stands out as a potent tool for navigating through

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: An illustrative example showing the difficult-to-define constraints using three districting plans for the Atlanta Police Department (APD). Gray lines represent the basic geographical units patrolled by the police, red lines outline the districting plans, and dashed blue lines highlight the changes made to the pre-2019 plan. (a) and (b) are feasible plans implemented by the APD pre and post 2019. (c) appears to be a feasible plan but was ultimately rejected by the APD because it overlooked traffic constraints and inadvertently cut off access to some highways with its zone boundaries.

these intricate decision spaces, providing a means to optimize outcomes without necessitating a detailed understanding of the underlying objective function’s analytical properties.

However, the broader application prospects of BBO in public policymaking are hampered due to two major hurdles: (1) Defining the feasible region or setting clear constraints for decisions is inherently complicated for real human systems. Policymakers often encounter a myriad of both explicit and implicit rules when making an optimal decision, adding a significant layer of complexity to the optimization process. For instance, in police districting, police departments often organize their patrol forces by dividing the geographical region of a city into multiple patrol areas called zones. Their goal is to search for the optimal districting plan that minimizes the workload variance across zones (Larson 1974; Larson and Odoni 1981; Zhu et al. 2020; Zhu, Wang, and Xie 2022). As shown in Figure 1, a well-conceived plan necessitates each zone to adhere to certain shape constraints (e.g., contiguity and compactness) that are analytically challenging to formulate (Shirabe 2009), while

also taking socio-economic or political considerations into account (e.g., ensuring fair access to public facilities). This creates a web of *implicit constraints* (Choi et al. 2000; Audet, Caporossi, and Jacquet 2020) that are elusive to define clearly and have a complex high-dimensional structure (e.g., manifold shape), making the assessment of feasible region nearly as expensive as evaluating the objective itself (Gardner et al. 2014). (2) The decisions for public policy are usually high-dimensional, presenting a significant computational hurdle to utilizing traditional BBO methods (Luong et al. 2019; Wan et al. 2021; Binois and Wycoff 2022). For example, the redistricting problem can be formulated by mixed-integer programming, grappling with hundreds or even thousands of decision variables even for medium-sized service systems (Zhu et al. 2020; Zhu, Wang, and Xie 2022).

Despite the difficulty in formulating constraints for decisions in public policymaking, the rich repository of historical decisions adopted by the practitioners, combined with the increasingly easier access to human systems (van Veenstra and Kotterink 2017; Yu et al. 2021), offer a wealth of decision samples. Collecting these samples might involve seeking guidance from official public entities on decision feasibility or generating decisions grounded in domain expertise, bypassing the need to understand the explicit form of the constraints. These readily available decisions harbor implicit knowledge that adeptly captures the dynamics of implicit constraints, providing a unique opportunity to skillfully address these issues. This inspires us to develop a *conditional generative representation model* that maps the feasible region in the original space to a lower-dimensional latent space, which encapsulates the key pattern of these implicit constraints. As a result, the majority of existing BBO methods can be directly applied to solve the original optimization problem in this latent space without constraints.

In this paper, we aim to solve Implicit-Constrained Black-Box Optimization (ICBBO) problems. In these problems, *the constraints are not analytically defined, however the feasibility of a given decision can be easily verified*. We introduce a new approach called Conditional And Generative Black-box Optimization (CageBO). Using a set of labeled decisions as feasible or infeasible, we first construct a conditional generative representation model based on the conditional variational autoencoder (CVAE) (Sohn, Lee, and Yan 2015) to explore and generate new feasible decisions within the intricate feasible region. To mitigate the impact of potentially poorly generated decisions, we adopt a post-decoding process that aligns these decisions closer to the nearest feasible ones. Furthermore, we incorporate this conditional generative representation model as an additional surrogate into our black-box optimization algorithm, allowing for a two-way mapping between the feasible region in the original space and an unconstrained latent space. As a result, while objective function is evaluated in the original space, the black box optimization algorithm is performed in the latent space without constraints. We prove that our proposed algorithm can achieve a no-regret upper bound with a judiciously chosen number of observations in the original space. Finally, we validate our approach with numerical experiments on both synthetic and real-world datasets, includ-

ing applying our method to address large-scale redistricting challenges within police operation systems in Atlanta, Georgia, demonstrating its significant empirical performance and efficiency over existing methods.

**Contributions.** Our contributions are summarized as:

1. We formulate a novel class of optimization problems in public policy, termed Implicit Constrained Black-Box Optimization (ICBBO), and develop the CageBO algorithm that effectively tackles high-dimensional ICBBO problems.
2. We introduce a conditional generative representation model that constructs a lower-dimensional, constraint-free latent space, enabling BBO algorithms to efficiently explore and generate feasible candidate solutions. This model supports the development of innovative public policy decisions by facilitating the exploration of new and effective policy options.
3. By establishing a no-regret expected cumulative regret bound, we demonstrate that our CageBO algorithm can consistently identify the global optimal solution, ensuring the development of near-optimal public policy decisions.
4. We apply our CageBO algorithm to the Atlanta police redistricting problem using both synthetic and real datasets. Our results demonstrate its superior performance compared to baseline methods, highlighting its potential to significantly enhance efficiency in real-world public policy applications.

**Related work.** Black-box optimization, a.k.a. zeroth-order optimization or derivative-free optimization, is a long-standing challenging problem in optimization and machine learning. Existing work either assumes the underlying objective function is drawn from some Gaussian process (Williams and Rasmussen 2006) or some parametric function class (Dai et al. 2022; Liu and Wang 2023). The former one is usually known as Bayesian optimization (BO), with the Gaussian process serving as the predominant surrogate model. BO has been widely used in many applications, including but not limited to neural network hyperparameter tuning (Kandasamy et al. 2020; Turner et al. 2021), material design (Ueno et al. 2016; Zhang, Apley, and Chen 2020), chemical reactions (Guo, Ranković, and Schwaller 2023), and public policy (Xing and Hua 2022).

In numerous real-world problems, optimization is subject to various types of constraints. Eriksson proposes scalable BO with known constraints in high dimensions (Eriksson and Poloczek 2021). Letham explores BO in experiments featuring noisy constraints (Letham et al. 2019). Gelbart pioneered the concept of BO with unknown constraints (Gelbart, Snoek, and Adams 2014), later enhanced by Aria through the ADMM framework (Ariafar et al. 2019). Their constraints are unknown due to uncertainty but can be evaluated using probabilistic models. In addition, Choi and Audet study the unrelaxable hidden constraints in a similar way (Choi et al. 2000; Audet, Caporossi, and Jacquet 2020), where the feasibility of a decision can be evaluated by another black-box function. In contrast, the implicit constraints are often unknown due to the lack of analytical formulations in many public policy-making problems.

Building on latent space methodologies, Varol presented a constrained latent variable model integrating prior knowledge (Varol et al. 2012). Eissman (Eissman et al. 2018) presents a VAE-guided Bayesian optimization algorithm with attribute adjustment. Deshwal and Doppa focus on combining latent space and structured kernels over combinatorial spaces (Deshwal and Doppa 2021). Maus further investigates structured inputs in local spaces (Maus et al. 2022), and Antonova introduces dynamic compression within variational contexts (Antonova et al. 2020). However, it’s worth noting that none of these studies consider any types of constraints in their methodologies.

In public policy-making, implicit constraints are often a critical consideration in queueing service models. The pioneering work on the hypercube queueing model, introduced in Larson’s paper (Larson 1974) and further developed in his book (Larson and Odoni 1981), established a foundational framework for analyzing spatial queueing systems, particularly in the context of emergency services like police and ambulance operations. This model is frequently employed as a black-box performance measure for public resource deployment (de Souza et al. 2015; Ansari, McLay, and Mayorga 2017; Zhu et al. 2020). However, real-world service systems are often subject to various constraints on feasible policies. These constraints, such as fairness (Argyris, Karsu, and Yavuz 2022) and continuity (Zhu, Wang, and Xie 2022), are typically implicit and lack explicit analytical formulations, posing challenges for effective optimization.

## Preliminaries

**Problem setup.** We consider a decision space denoted by  $\mathcal{X} \subseteq [0, 1]^d$  where 1 can be replaced with any universal constant w.l.o.g., which represents a specific region of a  $d$ -dimensional real space. Suppose there exists a black-box objective function,  $f : \mathcal{X} \mapsto \mathbb{R}$ , that can be evaluated, albeit at a substantial cost. Assume we can obtain a noisy observation of  $f(x)$ , denoted as  $\hat{f}(x) = f(x) + \epsilon$ , where  $\epsilon$  follows a  $\sigma$ -sub-Gaussian noise distribution. The goal is to solve the following optimization problem:

$$\min_{x \in \mathcal{X}} f(x), \quad \text{s.t. } x \in \mathcal{S},$$

where  $\mathcal{S} \subseteq \mathcal{X}$  represents the feasible region, defined by a set of implicit constraints.

Given the analytical expressions of the implicit constraints are not directly accessible, explicitly formulating these constraints is not feasible. However, they can still be evaluated through a feasibility oracle  $h(\cdot) : \mathcal{X} \mapsto \{0, 1\}$ , where a value of 1 indicates feasible, and 0 indicates infeasible. Now suppose we have access to a human system that provides labeled decisions. Denote a set of labeled decisions by  $\mathcal{D} = \{(x_i, c_i)\}$ ,  $i \in [n]$ , where  $x_i \in \mathcal{X}$  represents the  $i$ th decision and  $c_i \in \{0, 1\}$  represents its feasibility. Assume  $\mathcal{D}$  has a good coverage of the decision space of interest. In practice, decisions can be derived from consultations within human systems or crafted using domain expertise, with the feasibility oracle effectively acting as a surrogate for a policymaker. For example, new feasible districting plans in Figure 1 can be created by first randomly altering the assign-

ments of border regions and then checking their feasibility through police consultations.

**Bayesian optimization.** The BO algorithms prove especially valuable in scenarios where the evaluation of the objective function is costly or time-consuming, or when the gradient is unavailable. This approach revolves around constructing a *surrogate model* of the objective function and subsequently employing an *acquisition function* based on this surrogate model to determine the next solution for evaluation. For the minimization problem, a popular choice of surrogate model is the Gaussian process with the lower confidence bound (LCB) (Srinivas et al. 2009) serving as the acquisition function.

The Gaussian process (GP) in the space  $\mathcal{X}$ , denoted by  $\text{GP}(\mu, k; \mathcal{X})$ , is specified by a mean function  $\mu(x)$  and a kernel function  $k(x, x')$ , which indicates the covariance between the two arbitrary decisions  $x$  and  $x'$ . The GP captures the joint distribution of all the evaluated decisions and their observed objective function values. We reference the standard normal distribution with zero mean and an identity matrix  $I$  as its variance by  $\mathcal{N}(0, I)$ , and let  $Y = \{\hat{f}(x) \mid x \in X\}$  represent the corresponding set of objective function values. For a new decision  $\tilde{x}$ , the joint distribution of  $Y$  and its objective function value  $\tilde{y}$  of  $\tilde{x}$  is

$$\begin{bmatrix} Y \\ \tilde{y} \end{bmatrix} \sim \mathcal{N} \left( \mu \left( \begin{bmatrix} X \\ \tilde{x} \end{bmatrix} \right), \begin{bmatrix} K(X, X) + \sigma^2 I & K(X, \tilde{x}) \\ K(X, \tilde{x})^\top & k(\tilde{x}, \tilde{x}) \end{bmatrix} \right),$$

where  $\sigma^2$  represent the variance of the observed noise  $\epsilon$ . Here  $K(X, X) = (k(x, x'))_{x, x' \in X}$  denotes the covariance matrix between the previously evaluated decisions and  $K(X, \tilde{x}) = (k(x, \tilde{x}))_{x \in X}$  denotes the covariance vector between the previously evaluated and the new decisions.

## Proposed Method

The main idea of our method is to learn a conditional generative representation of feasible decisions to effectively overcome the complications posed by implicit constraints. Then we perform Bayesian optimization within this low-dimensional representation space rather than the constrained original decision space. The latent space  $\mathcal{Z} \subseteq \mathbb{R}^d$  is learned using a conditional variational autoencoder (CVAE), which leverages the set of labeled decisions  $\mathcal{D}$  as the training data.

To be specific, we train the CVAE model by maximizing  $\mathcal{L}_{\text{ELBO}}$  on the training data  $\mathcal{D}$ . This CVAE model enables a two-way mapping of decisions between the original and latent spaces, through an *encoder* $_{\phi}$  from the learned conditional distribution  $q_{\phi}(z|x, c)$  and a *decoder* $_{\theta}$  from the learned conditional posterior  $p_{\theta}(x|z, c)$ . Given an initial set of decisions  $X_0$  randomly sampled from the feasible set  $\mathcal{D}_f$ , we first evaluate their objective values denoted as  $Y_0 = \{\hat{f}(x) \mid x \in X_0\}$ . These decisions are then encoded to the latent space, represented by  $Z_0 \subset \mathcal{Z}$ . As illustrated by Figure 2, for each iteration  $t$ , our algorithm is performed as follows: (1) Train a surrogate model  $\text{GP}(\mu, k; \mathcal{Z})$  using the current latent decisions  $Z_{t-1}$  and their observed values  $Y_{t-1}$ . (2) Identify the next latent decision candidate  $z_t$  by searching on the feasible latent space  $\hat{\mathcal{Z}}$ , defined as the support of the latent variable distribution conditioned on feasibility




Figure 2: An illustration of the CageBO algorithm. Red dots represent observed feasible decisions, and grey dots denote observed infeasible decisions. The red dotted circle illustrates the complex feasible region that is not directly accessible. For each iteration, a new decision  $z_t$  is then chosen in the feasible latent space  $\hat{Z}$  by minimizing the lower confidence bound (LCB). This new decision  $z_t$  is mapped back to the original space as  $x_t$  using  $\text{decoder}_\theta$ . If the CVAE model is well-trained, the newly generated  $x_t$  is highly likely to reside within the feasible set  $\mathcal{S}$ . In case  $x_t \notin \mathcal{S}$ , the post-decoding process  $\psi_{\mathcal{D}_f}$  will adjust it to the nearest observed feasible decision.

$c = 1$ , and select the one exhibiting the lowest LCB value. (3) Decode the latent decision candidate  $z_t$  to the original space, yielding a new decision  $x_t$ . (4) Assess the feasibility of  $x_t$  via the feasibility oracle  $h$ . If  $x_t$  is feasible, evaluate its objective value and include it in the feasible set  $\mathcal{D}_f$ . Otherwise, apply the post-decoding process  $\psi_{\mathcal{D}_f}$  before evaluating the objective value. (5) Update the observations  $Z_t$  and  $Y_t$  accordingly. The CageBO algorithm iterates a total of  $T$  times. The proposed method is summarized in Algorithm 1. In the remainder of this section, we explain each component of our proposed method at length.

### Conditional Generative Representation

To address the challenge of implicit constraints in ICBBO problems, we introduce a conditional generative representation model based on a CVAE in our framework. The rationale behind our approach is threefold: (1) To encode the original decision space with implicit constraints into a compact, continuous, and constraint-free latent space. (2) To condense the dimensionality of the original problem, making BBO more efficient in the latent space. (3) To actively search for solutions with a high likelihood of being feasible. As illustrated by Figure 3, we observe that BBO can navigate within a feasible latent space, which offers a simpler structure in the objective function, facilitated by the conditional generative representation model.

Suppose there exists a joint distribution between the decision variable  $x \in \mathcal{X}$ , which can be defined in a high-dimensional hybrid space of discrete and continuous variables, and a continuous latent variable  $z \in \mathcal{Z} \subseteq \mathbb{R}^d$ , conditioned on the feasibility criterion  $c \in \{0, 1\}$ . We model the conditional distributions  $q_\phi(z|x, c)$ ,  $p_\theta(x|z, c)$ , and the conditional prior  $p(z|c)$  through the use of neural networks

---

### Algorithm 1: CageBO

---

**Input:** Labeled decisions  $\mathcal{D}$ ; Objective function  $\hat{f}$ ; Feasibility oracle  $h$ ; Hyper-parameters  $T$ .  
 Train a CVAE( $p_\theta, q_\phi$ ) by maximizing  $\mathcal{L}_{\text{ELBO}}$  in  $\mathcal{D}$ ;  
 Feasible set  $\mathcal{D}_f = \{x_i \in \mathcal{X} \mid (x_i, c_i) \in \mathcal{D} \text{ and } c_i = 1\}$ ;  
 Initialize  $X_0 \subseteq \mathcal{D}_f, Y_0 = \{\hat{f}(x) \mid x \in X_0\}$ ;  
 $Z_0 \leftarrow \text{encoder}_\phi(X_0)$ ; // Encoding  
**for**  $t = 1$  **to**  $T$  **do**  
   Train GP( $\mu, k; \mathcal{Z}$ ) using  $(Z_{t-1}, Y_{t-1})$ ;  
    $\hat{Z} = \text{supp}(q_\phi(\cdot|c = 1))$ ;  
    $z_t \in \text{argmin}_{z \in \hat{Z}} \text{LCB}(z, \text{GP})$ ;  
    $x_t \leftarrow \text{decoder}_\theta(z_t)$ ; // Decoding  
   **if**  $h(x_t) = 1$  **then**  
      $y_t = \hat{f}(x_t), \mathcal{D}_f \leftarrow \mathcal{D}_f \cup \{x_t\}$ ;  
   **else**  
      $y_t = \hat{f}(\psi_{\mathcal{D}_f}(x_t))$ ; // Post-decoding  
   **end if**  
    $Z_t \leftarrow Z_{t-1} \cup \{z_t\}, Y_t \leftarrow Y_{t-1} \cup \{y_t\}$ ;  
**end for**  
**return**  $x^* \in \text{argmin}_{x \in X_T} \hat{f}(x)$

---

as outlined in (Pinheiro Cinelli et al. 2021). Here,  $\phi$  and  $\theta$  represent the neural network weights associated with each respective distribution. Since it is intractable to directly optimize the marginal maximum likelihood of  $x$ , the evidence lower bound (ELBO) (Jordan et al. 1999) of the log-likelihood is derived as follows

$$\mathcal{L}_{\text{ELBO}} = w(c) \mathbb{E}_{q_\phi(z|x, c)} [\log p_\theta(x|z, c)] - \eta \text{D}_{\text{KL}}(q_\phi(z|x, c) \| p(z|c)), \quad (1)$$

where  $w(c)$  denotes the weighting function associated with the feasibility  $c$ , and  $\eta$  is a hyperparameter controlling the penalty ratio. Here we assume the conditional prior of  $z$  follows a standard Gaussian distribution  $\mathcal{N}(0, I)$ . The first term in (1) can be considered as the reconstruction error between the input and reconstructed decisions, and the second term is the Kullback–Leibler (KL) divergence between the conditional prior of the latent variable  $z$  and the learned posterior  $q_\phi(z|x, c)$ .

**Post-decoding process.** To ensure the decoded decision  $x \in \mathcal{X}$  is subject to the implicit constraints, we introduce a post-decoding process in addition to the decoder, denoted by  $\psi_{\mathcal{S}} : \mathcal{X} \mapsto \mathcal{S}$ . This function projects any given decision  $x \in \mathcal{X}$  to the closest feasible decision  $\hat{x} \in \mathcal{S}$ . However, the presence of implicit constraints prevents us from achieving an exact projection. As a workaround, we search within the observed feasible set  $\mathcal{D}_f = \{x_i \in \mathcal{X} \mid (x_i, c_i) \in \mathcal{D} \text{ and } c_i = 1\}$ , rather than the unattainable feasible region  $\mathcal{S}$ , and find a feasible decision  $\hat{x} \in \mathcal{D}_f$  that is the closest to the decoded decision  $x$  as an approximate. The distance between any two decisions is measured by the Euclidean norm  $\|\cdot\|_2$ . Formally,

$$\psi_{\mathcal{D}_f}(x) = \text{argmin}_{\hat{x} \in \mathcal{D}_f} \|\hat{x} - x\|_2. \quad (2)$$

The post-decoding process is initiated only for decoder-generated decisions  $x$  that are infeasible. Feasible decisions

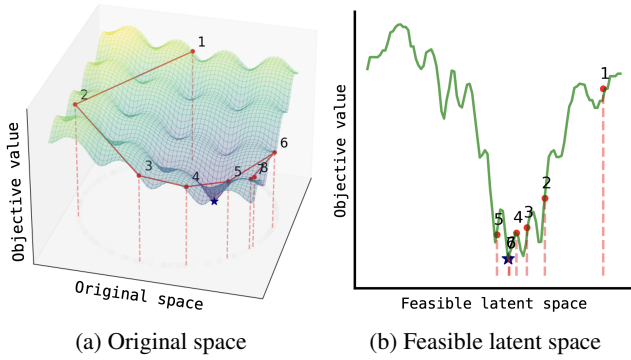


Figure 3: Solution paths of the same optimization problem suggested by our CageBO algorithm in the (a) original and (b) latent spaces, respectively. The objective in this illustrative example is to minimize the Ackley function subject to an "implicit" constraint (a circle) in a 2-dimensional decision space.

are directly incorporated into the observed feasible set  $\mathcal{D}_f$ . Through iteratively expanding this observed feasible set, we enhance the accuracy of this process by improving the coverage of  $\mathcal{D}_f$  on the underlying feasible region  $\mathcal{S}$ . This is further detailed in the ablation study in the appendix <sup>1</sup>.

### Surrogate Model

Now we define an *indirect objective function*  $g(\cdot)$  which maps from  $\mathcal{Z} \mapsto \mathbb{R}$ :

$$g(z) = \begin{cases} f(\text{decoder}_\theta(z)), & \text{if } h(\text{decoder}_\theta(z)) = 1, \\ f(\psi_{\mathcal{D}_f}(\text{decoder}_\theta(z))), & \text{if } h(\text{decoder}_\theta(z)) = 0. \end{cases} \quad (3)$$

The indirect objective function measures the objective value of the latent variable via the decoding and the potential post-decoding process. Given that the objective function  $f$  is inherently a black-box function, it follows that the indirect objective function  $g$  is also a black-box function.

We use a GP as our surrogate model of the indirect objective function  $g$ , denoted by  $\text{GP}(\mu, k; \mathcal{Z})$ . In our problem, the mean function can be written as  $\mu(z) = \mathbb{E}[g(z)]$ . In addition, we adopt the Matérn kernel (Seeger 2004) as the kernel function  $k(z, z') = \mathbb{E}[(g(z) - \mu(z))(g(z') - \mu(z'))]$ , which is widely-used in BO literature. The main advantage of the GP as a surrogate model is that it can produce estimates of the mean evaluation and variance of a new latent variable, which can be used to model uncertainty and confidence levels for the acquisition function described in the following. Note that the latent variable  $z$  is assumed to follow a Gaussian prior in the latent decision model, which aligns with the assumption of the GP model that the observed latent variables  $\mathcal{Z}$  follow the multivariate Gaussian distribution.

**Acquisition function.** In the BO methods, the acquisition function is used to suggest the next evaluating candidate. Our approach adopts the lower confidence bound (LCB) as

the acquisition function to choose the next latent variable  $z$  candidate to be decoded and evaluated. This function contains both the mean  $\mu(z)$  of the GP as the explicit exploitation term and the standard deviation  $\sigma(z)$  of the GP as the exploration term:

$$\text{LCB}(z, \text{GP}) = \mu(z) - \sqrt{\beta}\sigma(z), \quad (4)$$

where  $\beta$  is a trade-off parameter.

To identify new latent decisions for evaluation, our method draws numerous independent samples from the conditional posterior  $q_\phi(\cdot|x, c = 1)$  among observed feasible decisions  $x \in \mathcal{D}_f$ . The sample with the lowest LCB value, denoted as  $z_t$ , is selected for decoding and subsequent evaluation using the objective function  $\hat{f}$  in the  $t$ -th iteration.

### Theoretical Analysis

We provide theoretical analysis for Algorithm 1. Let  $\text{enc} : \mathcal{X} \mapsto \mathcal{Z}$  denote the encoder function,  $\text{dec} : \mathcal{Z} \mapsto \mathcal{X}$  denote the decoder function, and  $g(z) = f(\psi_{\mathcal{D}_f}(\text{dec}(z))) : \mathcal{Z} \mapsto \mathbb{R}$  denotes the objective function w.r.t.  $\mathcal{Z}$  where  $\psi_{\mathcal{D}_f}$  is the post-decoder. Note even if post-decoder is not needed,  $\psi_{\mathcal{D}_f}(\text{dec}(z)) = \text{dec}(z)$ . Following the existing work in Gaussian process bandit optimization, we utilize cumulative regret and its expected version to evaluate the performance of our algorithm which are defined as follows.

$$R_T = \sum_{i=1}^T f(x_i) - f(x_*), \mathbb{E}[R_T] = \sum_{i=1}^T \mathbb{E}[f(x_i) - f(x_*)], \quad (5)$$

where  $x_* = \text{argmin}_{x \in \mathcal{X}} f(x)$  and the expectation is taken over all randomness, including random noise and random sampling over observations.

We assume the distance between any two points in  $\mathcal{Z}$  can be upper bounded by their distance in  $\mathcal{X}$ , i.e.,  $\forall x, x' \in \mathcal{X}, \|\text{enc}(x) - \text{enc}(x')\|_2 \leq C_p \|x - x'\|_2$ . We further assume that function  $g : \mathcal{Z} \mapsto \mathbb{R}$  is drawn from a Gaussian Process prior and it is  $C_g$ -Lipschitz continuous, i.e.,  $\forall z, z' \in \mathcal{Z}, |g(z) - g(z')| \leq C_g \|z - z'\|_2$ . This assumption is the standard Gaussian process assumption for Bayesian optimization (Srinivas et al. 2010). Here  $C_p$  and  $C_g$  are universal constants. We use big  $\tilde{O}$  notation to omit any constant and logarithmic terms. Now we are ready to state our main theoretical result.

**Theorem 1.** After running  $T$  iterations, the expected cumulative regret of Algorithm 1 satisfies that

$$\mathbb{E}[R_T] = \tilde{O}(\sqrt{T\gamma_T} + \sqrt{d}(n + T)^{\frac{d}{d+1}}), \quad (6)$$

where  $\gamma_T$  is the maximum information gain, depending on choice of kernel used in algorithm and  $n$  is number of initial observation data points.

**Remark 2.** We present an upper bound of expected cumulative regret of our Algorithm 1. This is a no-regret algorithm since  $\lim_{T \rightarrow \infty} \mathbb{E}[R_T]/T = 0$  which means our algorithm is able to find the global optimal solution by expectation. The bound has two terms. The first term follows from GP-UCB (Srinivas et al. 2010) where the maximum information gain depends on the choice of kernel used in the algorithm. For

<sup>1</sup>Appendix is available on <https://arxiv.org/abs/2310.18449>. Code is available on <https://github.com/wenqian-xing/CageBO>.

linear kernel,  $\gamma_T = O(d \log T)$  and for squared exponential kernel,  $\gamma_T = O(\log T)^{d+1}$ . The second term is the regret term incurred by the post decoder  $\psi_{D_f}$  when needed, which is also sublinear in  $T$  if  $n$  is chosen no larger than  $T$ . Due to page limit, full proof of Theorem 1 is shown in the appendix.

## Experiments

### Experimental Setup

We compare our `CageBO` algorithm against three baseline approaches that can be used to address ICBBO problems. These baselines include (1) vanilla version of Bayesian optimization (BO) (Jones, Schonlau, and Welch 1998), (2) simulated annealing (SA) (Kirkpatrick, Gelatt Jr, and Vecchi 1983), (3) approximated Mixed Integer Linear Programming (MILP) (Zhu, Wang, and Xie 2022), (4) a VAE-guided Bayesian optimization (VAE-BO) (Eissman et al. 2018), and (5) a variant of constrained VAE-guided Bayesian optimization (CON-BO) (Griffiths and Hernández-Lobato 2020). We train the latent decision model in our framework with the Adam optimizer (Kingma and Ba 2015). Both `CageBO`, VAE-BO, and CON-BO are trained and performed under an identical environment. MILP is only compared in the redistricting problems because their implicit constraints can be approximated by a set of linear constraints, albeit with the trade-offs of adding auxiliary variables and incurring computational expenses. Each method is executed 10 times across all experiments to determine the 95% confidence interval of their results.

Experiments were conducted on a PC equipped with a 12-core CPU and 18 GB RAM. For synthetic experiments, `CageBO` is trained in 1,000 epochs with the learning rate of  $10^{-4}$ ,  $\eta = 0.1$ , and dimension of latent space  $d = 10$ . Hyperparameters:  $|X_0| = 10$  initial evaluation points,  $\beta = 1$  for the LCB, and  $T = 100$ . For real-world redistricting experiments, `CageBO` is trained in 1,000 epochs with the learning rate of  $10^{-4}$ ,  $\eta = 0.1$ , and the dimension of latent space  $d = 25$ . Hyperparameters:  $|X_0| = 5$  initial evaluation points,  $\beta = 1$  for the LCB, and total iterations  $T = 100$ .

### Synthetic Results

We consider minimizing (a) the 30-d Keane’s bump function (Keane 1994) and (b) the 30-d Michalewicz function (Molga and Smutnicki 2005), both of which are common test functions for constrained optimization. To obtain samples from implicit constraints, the conundrum we aim to address via our methodology, we first generate  $n = 2,000$  samples from the standard uniform distribution in the 10-dimensional latent space, then decode half of them through a randomly initialized decoder and mark those as feasible. We define the feasibility oracle to return 1 only if the input solution is matched with a feasible solution. We scale the samples accordingly so that the test functions can be evaluated under their standard domains, *i.e.*, the Keane’s bump function on  $[0, 10]^d$  and the Michalewicz function on  $[0, \pi]^d$ .

Figure 4(a)&(b) presents the synthetic results. It is evident that our method attains the lowest objective values consistently compared to other baseline methods. In Figure 4 (b),

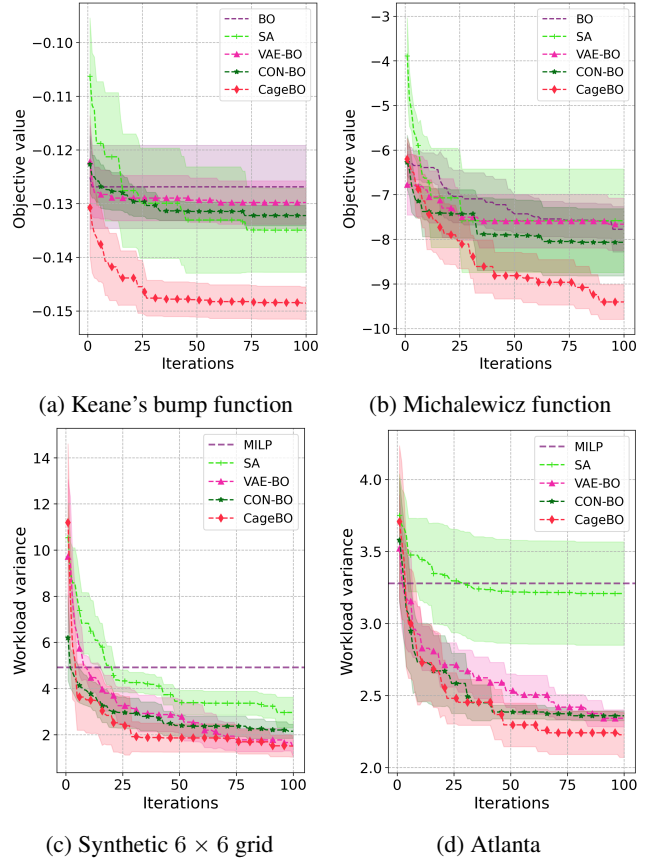


Figure 4: Comparisons of performance convergence with 95% confidence intervals for (a) 30-dimensional Keane’s bump function, (b) 30-dimensional Michalewicz function, (c) police redistricting problem in a synthetic  $6 \times 6$  grid, and (d) police redistricting problem in Atlanta, Georgia.

in particular, we observe that the integration of the conditional generative representation model into `CageBO` greatly enhances the BO’s performance. This is in stark contrast to the similar learning-based VAE-BO and CON-BO, which do not yield satisfactory outcomes.

### Case Study: Police Redistricting

One common application of high-dimensional ICBBO in public policymaking is redistricting. For example, in police redistricting problems, the goal is to distribute  $L$  police service regions across  $J$  distinct zones. Each service region is patrolled by a single police unit. While units within the same zone can assist each other, assistance across different zones is disallowed. The decision variables are defined by the assignment of a region to a zone, represented by matrix  $x \in \{0, 1\}^{L \times J}$ . Here, an entry  $x_{lj} = 1$  indicates region  $l$  is allocated to zone  $j$ , and  $x_{lj} = 0$  otherwise. A primal constraint is that each region should be assigned to only one zone. This means that for every region  $l$ ,  $\sum_{j \in [J]} x_{lj} = 1$ . The implicit constraint to consider is the contiguity constraint, which ensures that all regions within a specific zone

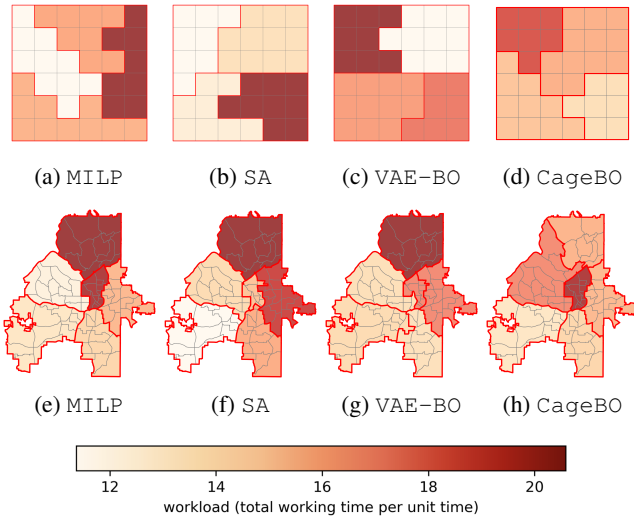


Figure 5: Districting plans for the police redistricting problem in a synthetic  $6 \times 6$  grid and a real-world scenario in Atlanta, Georgia. The workload in each region is indicated by the depth of the color. The *CageBO* plan achieves the lowest workload variance among all districts, indicating an optimal distribution of workload.

are adjacent. We define the feasibility oracle to return 1 only if both the primal and the implicit constraint are satisfied.

The objective of the police redistricting problem is to minimize the workload variance across all the districts. The workload for each district  $j$ , denoted as  $\rho_j$ , is computed using the following equation:

$$\rho_j = (\tau_j + 1/\mu)\lambda_j. \quad (7)$$

In this equation,  $\tau_j$  represents the average travel time within district  $j$ , which can be estimated using the hypercube queuing model (Larson 1974), which can be regarded as a costly black-box function. Assuming  $\mu$  represents the uniform service rate across all districts and  $\lambda_j$  denotes the arrival rate in district  $j$ . Essentially,  $\rho_j$  quantifies the cumulative working duration per unit time for police units in district  $j$ . For clarity, a value of  $\rho_j = 10$  implies that the combined working time of all police units in district  $j$  counts to 10 hours or minutes in every given hour or minute.

**Redistricting  $6 \times 6$  grid.** We first evaluate our algorithm using a synthetic scenario, consisting of a  $6 \times 6$  grid to be divided into 4 zones, and the decision variable  $x \in \{0, 1\}^{36 \times 4}$ . For each region  $l$ , the arrival rate  $\lambda_l$  is independently drawn from a standard uniform distribution. Moreover, all regions share an identical service rate of  $\mu = 1$ . We generate an initial data set  $\mathcal{D}$  by simulating  $n = 10,000$  labeled decisions.

Figure 4(c) demonstrates that our method notably surpasses other baseline methods regarding objective value and convergence speed. Additionally, in Figure 5, we show the optimal districting plans derived from our algorithm alongside those from the baseline methods. It is evident that our approach produces a more balanced plan compared to the other methods.

**Atlanta police redistricting.** For the police zone design in Atlanta, there are 78 police service regions that need to be divided into 6 zones, with the decision variable represented as  $x \in \{0, 1\}^{78 \times 6}$ . The redistricting of Atlanta’s police zones is constrained by several potential implicit factors. These include contiguity and compactness, as well as other practical constraints that cannot be explicitly defined. One such implicit constraint is the need for changes to the existing plan should be taken in certain local areas. A drastic design change is undesirable because: (1) A large-scale operational change will result in high implementation costs. (2) A radical design change will usually face significant uncertainties and unpredictable risks in future operations. The arrival rate  $\lambda_l$  and the service rate  $\mu$  are estimated using historical 911-calls-for-service data collected in the years 2021 and 2022 (Zhu and Xie 2019, 2022; Zhu, Wang, and Xie 2022). We generate an initial data set  $\mathcal{D}$  by simulating  $n = 10,000$  labeled decisions which are the neighbors of the existing plan with changes in certain local areas.

Figure 4(d) displays the convergence of our algorithm in comparison to other baseline methods, showing a consistent reduction in workload variance. Notably, the redistricting plan created by our *CageBO* algorithm most closely aligns with the actual heuristic policy applied by policymakers, wherein the zone with the highest arrival rate is assigned the largest workload, and the workloads of other zones are well-balanced. Further discussion on fair districting can be found in the appendix. As shown in Figure 5, the plan produced by our *CageBO* algorithm achieves the lowest workload variance, surpassing the baseline algorithms. The superior performance of our *CageBO* algorithm highlights its capacity to offer critical managerial insights to policymakers, thereby facilitating informed decision-making in real-world applications.

## Conclusion

In this paper, we introduce a new category of optimization problems in public policy, termed Implicit-Constrained Black-Box Optimization (ICBBO), where constraints are implicit but can be easily verified. This problem is well motivated by real-world police redistricting problems where districting plans are constrained to socio-economic or political considerations but feasible plans can be easily verified by policymakers. This new ICBBO problem poses more challenges than conventional black-box optimization where existing work cannot be readily applied.

To address this problem, we develop the *CageBO* algorithm and its core idea is to learn a conditional generative representation of feasible decisions to effectively overcome the complications posed by implicit constraints. In this way, our method can not only learn a good representation of the original space but also generate feasible samples that can be used later. Our theoretical analysis confirms that *CageBO* can, on expectation, identify the global optimal solution. This approach was empirically validated through a case study on police redistricting in Atlanta, Georgia, where our algorithm demonstrated superior performance in optimizing districting plans.

## Acknowledgments

This work was partially done when CL was at the University of Chicago. This work was partially supported by the UAI-bany Computer Science Department startup funding. The authors would like to thank the AAAI-AISI reviewers and area chair for helpful comments.

## References

- Ansari, S.; McLay, L. A.; and Mayorga, M. E. 2017. A maximum expected covering problem for district design. *Transportation Science*, 51(1): 376–390.
- Antonova, R.; Rai, A.; Li, T.; and Kragic, D. 2020. Bayesian Optimization in Variational Latent Spaces with Dynamic Compression. In Kaelbling, L. P.; Kragic, D.; and Sugiura, K., eds., *Proceedings of the Conference on Robot Learning*, volume 100 of *Proceedings of Machine Learning Research*, 456–465. PMLR.
- Argyris, N.; Karsu, Ö.; and Yavuz, M. 2022. Fair resource allocation: Using welfare-based dominance constraints. *European journal of operational research*, 297(2): 560–578.
- Ariaifar, S.; Coll-Font, J.; Brooks, D.; and Dy, J. 2019. ADMMBO: Bayesian Optimization with Unknown Constraints using ADMM. *Journal of Machine Learning Research*, 20(123): 1–26.
- Audet, C.; Caporossi, G.; and Jacquet, S. 2020. Binary, unrelaxable and hidden constraints in blackbox optimization. *Operations Research Letters*, 48(4): 467–471.
- Binois, M.; and Wycoff, N. 2022. A survey on high-dimensional Gaussian process modeling with application to Bayesian optimization. *ACM Transactions on Evolutionary Learning and Optimization*, 2(2): 1–26.
- Chandak, A.; Dey, D.; Mukhoty, B.; and Kar, P. 2020. Epidemiologically and socio-economically optimal policies via Bayesian optimization. *Transactions of the Indian National Academy of Engineering*, 5: 117–127.
- Choi, T. D.; Eslinger, O. J.; Kelley, C. T.; David, J. W.; and Etheridge, M. 2000. Optimization of automotive valve train components with implicit filtering. *Optimization and Engineering*, 1: 9–27.
- Dai, Z.; Shu, Y.; Low, B. K. H.; and Jaillet, P. 2022. Sample-then-optimize batch neural thompson sampling. *Advances in Neural Information Processing Systems*, 35: 23331–23344.
- de Souza, R. M.; Morabito, R.; Chiyoshi, F. Y.; and Iannoni, A. P. 2015. Incorporating priorities for waiting customers in the hypercube queuing model with application to an emergency medical service system in Brazil. *European journal of operational research*, 242(1): 274–285.
- Deshwal, A.; and Doppa, J. 2021. Combining latent space and structured kernels for bayesian optimization over combinatorial spaces. *Advances in Neural Information Processing Systems*, 34: 8185–8200.
- Eissman, S.; Levy, D.; Shu, R.; Bartzsch, S.; and Ermon, S. 2018. Bayesian optimization and attribute adjustment. In *Proc. 34th Conference on Uncertainty in Artificial Intelligence*.
- Eriksson, D.; and Poloczek, M. 2021. Scalable Constrained Bayesian Optimization. In Banerjee, A.; and Fukumizu, K., eds., *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, 730–738. PMLR.
- Gardner, J. R.; Kusner, M. J.; Xu, Z. E.; Weinberger, K. Q.; and Cunningham, J. P. 2014. Bayesian optimization with inequality constraints. In *ICML*, volume 2014, 937–945.
- Gelbart, M. A.; Snoek, J.; and Adams, R. P. 2014. Bayesian optimization with unknown constraints. In *30th Conference on Uncertainty in Artificial Intelligence, UAI 2014*, 250–259.
- Griffiths, R.-R.; and Hernández-Lobato, J. M. 2020. Constrained Bayesian optimization for automatic chemical design using variational autoencoders. *Chemical science*, 11(2): 577–586.
- Guo, J.; Ranković, B.; and Schwaller, P. 2023. Bayesian Optimization for Chemical Reactions. *Chimia*, 77(1/2): 31–31.
- Jones, D. R.; Schonlau, M.; and Welch, W. J. 1998. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13: 455–492.
- Jordan, M. I.; Ghahramani, Z.; Jaakkola, T. S.; and Saul, L. K. 1999. An introduction to variational methods for graphical models. *Machine learning*, 37: 183–233.
- Kandasamy, K.; Vysyaraju, K. R.; Neiswanger, W.; Paria, B.; Collins, C. R.; Schneider, J.; Póczos, B.; and Xing, E. P. 2020. Tuning hyperparameters without grad students: Scalable and robust bayesian optimisation with dragonfly. *The Journal of Machine Learning Research*, 21(1): 3098–3124.
- Keane, A. 1994. Experiences With Optimizers In Structural Design. In *In Proceedings of the conference on adaptive computing in engineering design and control*, volume 94, 14–27.
- Kingma, D.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*. San Diego, CA, USA.
- Kirkpatrick, S.; Gelatt Jr, C. D.; and Vecchi, M. P. 1983. Optimization by simulated annealing. *science*, 220(4598): 671–680.
- Larson, R. C. 1974. A hypercube queuing model for facility location and redistricting in urban emergency services. *Computers & Operations Research*, 1(1): 67–95.
- Larson, R. C.; and Odoni, A. R. 1981. *Urban operations research*.
- Letham, B.; Karrer, B.; Ottoni, G.; and Bakshy, E. 2019. Constrained Bayesian Optimization with Noisy Experiments. *Bayesian Analysis*, 14(2): 495 – 519.
- Liu, C.; and Wang, Y.-X. 2023. Global Optimization with Parametric Function Approximation. In *International Conference on Machine Learning*, 22113–22136. PMLR.
- Luong, P.; Gupta, S.; Nguyen, D.; Rana, S.; and Venkatesh, S. 2019. Bayesian optimization with discrete variables. In *AI 2019: Advances in Artificial Intelligence: 32nd Australasian Joint Conference, Adelaide, SA, Australia, December 2–5, 2019, Proceedings 32*, 473–484. Springer.

- Maus, N.; Jones, H.; Moore, J.; Kusner, M. J.; Bradshaw, J.; and Gardner, J. 2022. Local latent space bayesian optimization over structured inputs. *Advances in Neural Information Processing Systems*, 35: 34505–34518.
- Molga, M.; and Smutnicki, C. 2005. Test functions for optimization needs. *Test functions for optimization needs*, 101: 48.
- Pardalos, P. M.; Rasskazova, V.; Vrahatis, M. N.; et al. 2021. *Black Box Optimization, Machine Learning, and No-Free Lunch Theorems*. Springer.
- Pinheiro Cinelli, L.; Araújo Marins, M.; Barros da Silva, E. A.; and Lima Netto, S. 2021. Variational autoencoder. In *Variational Methods for Machine Learning with Applications to Deep Networks*, 111–149. Springer.
- Seeger, M. 2004. Gaussian processes for machine learning. *International journal of neural systems*, 14(02): 69–106.
- Shirabe, T. 2009. Districting modeling with exact contiguity constraints. *Environment and Planning B: Planning and Design*, 36(6): 1053–1066.
- Sohn, K.; Lee, H.; and Yan, X. 2015. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28.
- Srinivas, N.; Krause, A.; Kakade, S.; and Seeger, M. 2010. Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design. In *Proceedings of the 27th International Conference on Machine Learning*, 1015–1022.
- Srinivas, N.; Krause, A.; Kakade, S. M.; and Seeger, M. 2009. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*.
- Turner, R.; Eriksson, D.; McCourt, M.; Kiili, J.; Laaksonen, E.; Xu, Z.; and Guyon, I. 2021. Bayesian optimization is superior to random search for machine learning hyperparameter tuning: Analysis of the black-box optimization challenge 2020. In *NeurIPS 2020 Competition and Demonstration Track*, 3–26. PMLR.
- Ueno, T.; Rhone, T. D.; Hou, Z.; Mizoguchi, T.; and Tsuda, K. 2016. COMBO: An efficient Bayesian optimization library for materials science. *Materials discovery*, 4: 18–21.
- van Veenstra, A. F.; and Kotterink, B. 2017. Data-driven policy making: The policy lab approach. In *Electronic Participation: 9th IFIP WG 8.5 International Conference, ePart 2017, St. Petersburg, Russia, September 4-7, 2017, Proceedings 9*, 100–111. Springer.
- Varol, A.; Salzmann, M.; Fua, P.; and Urtasun, R. 2012. A constrained latent variable model. In *2012 IEEE conference on computer vision and pattern recognition*, 2248–2255. Ieee.
- Wan, X.; Nguyen, V.; Ha, H.; Ru, B.; Lu, C.; and Osborne, M. A. 2021. Think global and act local: Bayesian optimization over high-dimensional categorical and mixed search spaces. *arXiv preprint arXiv:2102.07188*.
- Williams, C. K.; and Rasmussen, C. E. 2006. *Gaussian processes for machine learning*. MIT press Cambridge, MA.
- Xie, W.; Nie, W.; Saffari, P.; Robledo, L. F.; Descote, P.-Y.; and Jian, W. 2021. Landslide hazard assessment based on Bayesian optimization–support vector machine in Nanping City, China. *Natural Hazards*, 109(1): 931–948.
- Xing, W.; and Hua, C. 2022. Optimal Unit Locations in Emergency Service Systems with Bayesian Optimization. In *INFORMS International Conference on Service Science*, 439–452. Springer.
- Yu, S.; Qing, Q.; Zhang, C.; Shehzad, A.; Oatley, G.; and Xia, F. 2021. Data-driven decision-making in COVID-19 response: A survey. *IEEE Transactions on Computational Social Systems*, 8(4): 1016–1029.
- Zhang, Y.; Apley, D. W.; and Chen, W. 2020. Bayesian optimization for materials design with mixed quantitative and qualitative variables. *Scientific reports*, 10(1): 4924.
- Zhu, S.; Bukharin, A. W.; Lu, L.; Wang, H.; and Xie, Y. 2020. Data-driven optimization for police beat design in south fulton, georgia. *arXiv preprint arXiv:2004.09660*.
- Zhu, S.; Wang, H.; and Xie, Y. 2022. Data-Driven Optimization for Atlanta Police-Zone Design. *INFORMS Journal on Applied Analytics*, 52(5): 412–432.
- Zhu, S.; and Xie, Y. 2019. Crime event embedding with unsupervised feature selection. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 3922–3926. IEEE.
- Zhu, S.; and Xie, Y. 2022. Spatiotemporal-textual point processes for crime linkage detection. *The Annals of Applied Statistics*, 16(2): 1151–1170.