



# Bayesian optimization with approximate set kernels

Jungtaek Kim<sup>1</sup> · Michael McCourt<sup>2</sup> · Tackgeun You<sup>1</sup> · Saehoon Kim<sup>3</sup> · Seungjin Choi<sup>4</sup>

Received: 19 September 2020 / Revised: 15 January 2021 / Accepted: 23 January 2021 /

Published online: 22 March 2021

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2021

## Abstract

We propose a practical Bayesian optimization method over sets, to minimize a black-box function that takes a set as a single input. Because set inputs are permutation-invariant, traditional Gaussian process-based Bayesian optimization strategies which assume vector inputs can fall short. To address this, we develop a Bayesian optimization method with *set kernel* that is used to build surrogate functions. This kernel accumulates similarity over set elements to enforce permutation-invariance, but this comes at a greater computational cost. To reduce this burden, we propose two key components: (i) a more efficient approximate set kernel which is still positive-definite and is an unbiased estimator of the true set kernel with upper-bounded variance in terms of the number of subsamples, (ii) a constrained acquisition function optimization over sets, which uses symmetry of the feasible region that defines a set input. Finally, we present several numerical experiments which demonstrate that our method outperforms other methods.

**Keywords** Global optimization · Bayesian optimization · Set optimization

## 1 Introduction

Bayesian optimization is an effective method to optimize an expensive black-box function. It has proven useful in several applications, including hyperparameter optimization (Snoek et al. 2012; Hutter et al. 2011), neural architecture search (Zoph and Le 2017; Kandasamy et al. 2018), material design (Frazier and Wang 2016; Haghaniifar et al. 2019), and synthetic gene design (González et al. 2014). Classic Bayesian optimization assumes a search region  $\mathcal{X} \subset \mathbb{R}^d$  and a black-box function  $f$  evaluated in the presence of additive noise  $\epsilon$ , i.e.,  $y = f(\mathbf{x}) + \epsilon$  for  $\mathbf{x} \in \mathcal{X}$ .

---

Editors: Annalisa Appice, Sergio Escalera, Jose A. Gamez, Heike Trautmann

---

✉ Jungtaek Kim  
jtkim@postech.ac.kr

<sup>1</sup> POSTECH, Pohang, Republic of Korea

<sup>2</sup> SigOpt, an Intel company, San Francisco, USA

<sup>3</sup> Kakao Brain, Seongnam, Republic of Korea

<sup>4</sup> BARO AI, Seoul, Republic of Korea

Unlike this standard Bayesian optimization formulation, we assume that a search region is  $\mathcal{X}_{\text{set}} = \{\{\mathbf{x}_1, \dots, \mathbf{x}_m\} \mid \mathbf{x}_i \in \mathcal{X} \subset \mathbb{R}^d\}$  for a fixed positive integer  $m$ . Thus, for  $\mathbf{X} \in \mathcal{X}_{\text{set}}$ ,  $f$  would take in a *set* containing  $m$  elements, all of length  $d$ , and return a noisy function value  $y$ :

$$y = f(\mathbf{X}) + \epsilon. \quad (1)$$

Our motivating example comes from the soft  $k$ -means clustering algorithm over a data-set  $\mathcal{P} = \{\mathbf{p}_1, \dots, \mathbf{p}_N\}$ ; in particular, we aim to find the optimal initialization of such an algorithm. The objective function for this problem is a squared loss function which takes in the cluster initialization points  $\{\mathbf{x}_1, \dots, \mathbf{x}_k\}$  and returns the weighted distance between the points in  $\mathcal{P}$  and the converged cluster centers  $\{\mathbf{c}_1, \dots, \mathbf{c}_k\}$ . See Lloyd (1982) for more details.

Some previous research has attempted to build Gaussian process (GP) models on set data. Garnett et al. (2010) propose a method over discrete sets using stationary kernels over the first Wasserstein distance between two sets, though the power set of fixed discrete sets as domain space is not our interest. However, this method needs the complexity  $\mathcal{O}(n^2 m^3 d)$  to compute a covariance matrix with respect to  $n$  sets. Moreover, since it only considers stationary kernels, GP regression is restricted to a form that cannot express non-stationary models (Paciorek and Schervish 2004).

Therefore, we instead adapt and augment a strategy proposed in Gärtner et al. (2002) involving the creation of a specific *set kernel*. This set kernel uses a kernel defined on the elements  $\mathbf{x} \in \mathbb{R}^d$  of the sets to build up its own sense of covariance between sets. In turn, then, it can be directly used to build surrogate functions through GP regression, which can power the Bayesian optimization strategy, by Lemma 1.

A key contribution of this article is the development of a computationally efficient approximation to this set kernel. Given  $n$  total observed function values, the cost of constructing the matrix required for fitting the GP is  $\mathcal{O}(n^2 m^2 d)$  where  $m \geq n$  (see the complexity analysis in Sect. 3.3). The approximate set kernel proposed in this work uses random subsampling to reduce the computational cost to  $\mathcal{O}(n^2 L^2 d)$  for  $L < m$  while still producing an unbiased estimate of the expected value of the true kernel.

Another primary contribution is a constrained acquisition function optimization over set inputs. The next query set to observe is found by optimizing the acquisition function defined on a set  $\mathbf{X} \in \mathcal{X}_{\text{set}}$ . Using the symmetry of the space, this function can be efficiently optimized with a rejection sampling. Furthermore, we provide a theoretical analysis on cumulative regret bounds of our framework, which guarantees the convergence quality in terms of iterations.

## 2 Background

In this section, we briefly introduce previous studies, notations, and related work necessary to understand our algorithm.

### 2.1 Bayesian optimization

Bayesian optimization seeks to minimize an unknown function  $f$  which is expensive to evaluate,  $\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ , where  $\mathcal{X} \subset \mathbb{R}^d$  is a compact space. It is a sequential optimization strategy which, at each iteration, performs the following three computations:

1. Using the  $n$  data presently available,  $\{(\mathbf{x}_i, y_i)\}$  for  $i \in [n]$ , build a probabilistic surrogate model  $s_n$  meant to approximate  $f$ .
2. Using the surrogate model  $s_n$ , compute an acquisition function  $a_n$ , which represents the utility of next acquiring data at some new point  $\mathbf{x}$ .
3. Observe  $y_{n+1}$  from a true function  $f$  at the location  $\mathbf{x}_{n+1} = \arg \max_{\mathbf{x} \in \mathcal{X}} a_n(\mathbf{x})$ .

After exhausting a predefined budget  $T$ , Bayesian optimization returns the best point,  $\mathbf{x}^\dagger$ , that has the minimum observation. The benefit of this process is that the optimization of the expensive function  $f$  has been replaced by the optimization of much cheaper and better understood acquisition functions  $a_n$ .

In this paper, we use GP regression (Rasmussen and Williams 2006) to produce the surrogate function  $s_n$ ; from  $s_n$ , we use the Gaussian process upper confidence bound (GP-UCB) criterion (Srinivas et al. 2010):  $a_n(\mathbf{x}) = -\mu_n(\mathbf{x}) + \beta_n \sigma_n(\mathbf{x})$ , where  $\mu_n(\cdot)$  and  $\sigma_n^2(\cdot)$  are posterior mean and variance functions computed by  $s_n$ , and  $\beta_n$  is a trade-off hyperparameter for exploration and exploitation at iteration  $n$ . See Brochu et al. (2010), Shahriari et al. (2016), and Frazier (2018) for the details.

## 2.2 Set kernel

We introduce the notation required for performing kernel approximation of functions on sets. A set of  $m$  vectors is denoted as  $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ , where  $\mathbf{x}_i$  is in a compact space  $\mathcal{X} \subset \mathbb{R}^d$ . In a collection of  $n$  such sets (as will occur in the Bayesian optimization setting), the  $k$ th set would be denoted  $\mathbf{X}^{(k)} = \{\mathbf{x}_1^{(k)}, \dots, \mathbf{x}_m^{(k)}\}$ . Note that we are restricting all sets to be of the same size  $|\mathbf{X}^{(k)}| = m$ .<sup>1</sup>

To build a GP surrogate, we require a prior belief of the covariance between elements in  $\mathcal{X}_{\text{set}} = \{\{\mathbf{x}_1, \dots, \mathbf{x}_m\} \mid \mathbf{x}_i \in \mathcal{X} \subset \mathbb{R}^d\}$ . This belief is imposed in the form of a positive-definite kernel  $k_{\text{set}} : \mathcal{X}_{\text{set}} \times \mathcal{X}_{\text{set}} \rightarrow \mathbb{R}$ ; see Schölkopf and Smola (2002), Fasshauer and McCourt (2015) for more discussion on approximation with kernels. In addition to the symmetry  $k_{\text{set}}(\mathbf{X}^{(i)}, \mathbf{X}^{(j)}) = k_{\text{set}}(\mathbf{X}^{(j)}, \mathbf{X}^{(i)})$  required in standard kernel settings, kernels on sets require an additional property: the ordering of elements in  $\mathbf{X}$  should be immaterial (since sets have no inherent ordering).

Given an empirical approximation of the kernel mean  $\boldsymbol{\mu}_{\mathbf{X}} \approx |\mathbf{X}|^{-1} \sum_{i=1}^{|\mathbf{X}|} \phi(\mathbf{x}_i)$ , where  $\phi$  is a feature map  $\mathbb{R}^d \rightarrow \mathbb{R}^{d'}$  and  $d'$  is a dimensionality of projected space by  $\phi$ , a set kernel (Gärtner et al. 2002; Muandet et al. 2017) is defined as

$$k_{\text{set}}(\mathbf{X}^{(1)}, \mathbf{X}^{(2)}) = \langle \boldsymbol{\mu}_{\mathbf{X}^{(1)}}, \boldsymbol{\mu}_{\mathbf{X}^{(2)}} \rangle = \frac{1}{|\mathbf{X}^{(1)}| |\mathbf{X}^{(2)}|} \sum_{i=1}^{|\mathbf{X}^{(1)}|} \sum_{j=1}^{|\mathbf{X}^{(2)}|} k(\mathbf{x}_i^{(1)}, \mathbf{x}_j^{(2)}), \quad (2)$$

where  $k(\mathbf{x}_i, \mathbf{x}_j) = \langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle = \phi(\mathbf{x}_i)^\top \phi(\mathbf{x}_j)$ . Here,  $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$  is a positive-definite kernel defined to measure the covariance between the  $d$ -dimensional elements of the sets (e.g., a squared exponential or Matérn kernel). The kernel (2) arises when comparing a class of functions on different probability measures with the intent of understanding if the measures might be equal (Gretton et al. 2012).

<sup>1</sup> In principle, sets of varying size can be considered, but we restrict to same sized sets to simplify our analysis.

## 2.3 Related work

Although it has been raised in different interests, meta-learning approaches dealt with set inputs are promising in a machine learning community, because they can generalize distinct tasks with meta-learners (Edwards and Storkey 2017; Zaheer et al. 2017; Finn et al. 2017; Garnelo et al. 2018). In particular, they propose feed-forward neural networks which take permutation-invariant and variable-length inputs: they have the goal of obtaining features derived from the sets with which to input to a standard (meta-)learning routine. Since they consider modeling of set structure, they are related to our work, but they are interested in their own specific examples such as point cloud classification, few-shot learning, and image completion.

In Bayesian optimization literature, Garnett et al. (2010) suggest a method to find a set that produces a global minimum with respect to discrete sets, each of which is an element of power set of entire set. This approach solves the problem related to set structure using the first Wasserstein distance over sets. However, for the reason why the time complexity of the first Wasserstein distance is  $\mathcal{O}(n^2 m^3 d)$ , they assume a small cardinality of sets and discrete search space for the global optimization method. Furthermore, their method restricts the number of iterations for optimizing an acquisition function, since the number of iterations should increase exponentially due to the curse of dimensionality. This implies that finding the global optimum of acquisition function is hard to achieve.

Compared to Garnett et al. (2010), we consider continuous domain space which implies an acquired set is composed of any instances in a compact space  $\mathcal{X}$ . We thus freely use off-the-shelf global optimization method or local optimization method (Shahriari et al. 2016) with relatively large number of instances in sets. In addition, its structure of kernel is  $k_{\text{st}}(d(\mathbf{X}^{(1)}, \mathbf{X}^{(2)}))$  where  $k_{\text{st}}(\cdot)$  is a stationary kernel (Genton 2001) and  $d(\cdot, \cdot)$  is a distance function over two sets (e.g., in Garnett et al. 2010 the first Wasserstein distance). Using the method proposed in Sect. 3, a non-stationary kernel might be considered in modeling a surrogate function.

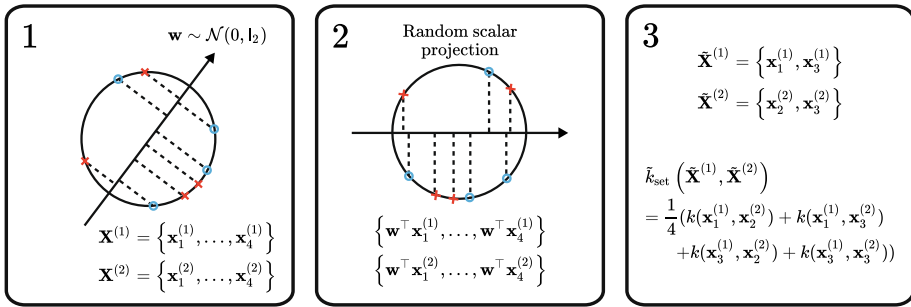
Recently, Buathong et al. (2020) solve a similar set optimization problem using Bayesian optimization with deep embedding kernels.<sup>2</sup> Compared to our method, it employs a kernel over RKHS embeddings as a kernel for set inputs, and shows its strict positive definiteness.

## 3 Proposed method

We first propose and analyze an approximation to the set kernel (2) for GP regression in this section. Then, we present a Bayesian optimization framework over sets, by introducing our Bayesian optimization with approximate set kernels and a constrained optimization method for finding the next set to evaluate.

In order for (2) to be a viable kernel of a GP regression, it must be positive-definite. To discuss this topic, we denote a list of  $n$  sets with the notation  $\mathfrak{X} = [\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(n)}] \in \mathcal{X}_{\text{set}}^n$ ; in this notation, the order of the entries matters.

<sup>2</sup> Although the work (Buathong et al. 2020) refers to our preliminary non-archival presentation (Kim et al. 2019), we mention (Buathong et al. 2020) here due to a close relationship with this work and its importance.



**Fig. 1** Illustration that shows how to select  $L$  instances from sets, which originally have  $m$  instances. In this example,  $m = 4$  and  $L = 2$ . (Phase 1) Two set inputs are projected onto a vector that is randomly drawn from the standard Gaussian distribution. The points that have same color belong to same set (e.g., blue and red). (Phase 2) The location of the projections onto the line determines the order of the instances. (Phase 3) Using the order of instances, two instances uniformly sampled are selected and they are used to compute the approximate set kernel value (Color figure online)

**Lemma 1** Suppose we have a list  $\mathfrak{X}$  which contains distinct sets  $\mathbf{X}^{(i)}$  for  $i \in [n]$ . We define the matrix  $\mathbf{K} \in \mathbb{R}^{n \times n}$  as

$$(\mathbf{K})_{ij} = k_{\text{set}}(\mathbf{X}^{(i)}, \mathbf{X}^{(j)}), \tag{3}$$

for  $k_{\text{set}}$  defined with a chosen inner kernel  $k$  as in (2). Then,  $\mathbf{K}$  is a symmetric positive-semidefinite matrix if  $k$  is a symmetric positive-definite kernel.

This proof appears in Haussler (1999, Lemma 1), and is also discussed in Gärtner et al. (2002).

### 3.1 Approximation of the set kernel

Computing (3) requires pairwise comparisons between all sets present in  $\mathfrak{X}$ , which has computational complexity  $\mathcal{O}(n^2 m^2 d)$ . To alleviate this cost, we propose to approximate (2) with

$$\tilde{k}_{\text{set}}(\mathbf{X}^{(1)}, \mathbf{X}^{(2)}; \pi, \mathbf{w}, L) = k_{\text{set}}(\tilde{\mathbf{X}}^{(1)}, \tilde{\mathbf{X}}^{(2)}), \tag{4}$$

where  $\pi : [m] \rightarrow [m]$ ,  $\mathbf{w} \in \mathbb{R}^d$  and  $L \in \mathbb{Z}_+$  and  $\tilde{\mathbf{X}}^{(i)}$  is a subset of  $\mathbf{X}^{(i)}$  which is defined by those three quantities (we omit explicitly listing them in  $\tilde{\mathbf{X}}^{(i)}$  to ease the notation).

The goal of the approximation strategy is to convert from  $\mathbf{X}^{(i)}$  (of size  $m$ ) to  $\tilde{\mathbf{X}}^{(i)}$  (of size  $L$ ) in a consistent fashion during all the  $\tilde{k}_{\text{set}}$  computations comprising  $\mathbf{K}$ . As shown in Fig. 1, we accomplish this in two steps:

1. Use a randomly generated vector  $\mathbf{w}$  to impose an (arbitrary) ordering of the elements of all sets  $\mathbf{X}^{(i)}$ , and
2. Randomly permute the indices  $[m]$  via a function  $\pi$ .

These random strategies are defined once before computing the  $\mathbf{K}$  matrix, and then used consistently throughout the entire computation.

---

**Algorithm 1** Forming the  $\tilde{K}$  approximation to  $K$

---

**Input:** A list of sets  $\mathcal{X} = [\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(n)}]$ ,  $L \in \mathbb{Z}_+$ .

**Output:** A matrix  $\tilde{K}$ .

- 1: Draw a random vector  $\mathbf{w}$  from  $\mathcal{N}(0, \mathbf{I}_d)$ .
- 2: Create a random permutation of the integers  $[m]$  with which to define  $\pi$ .
- 3: Using  $\mathbf{w}$ , assign the ordering of elements in  $\mathbf{X}^{(i)}$  according to (5).
- 4: Using  $\pi$ , determine the subsets  $\tilde{\mathbf{X}}^{(i)}$  as selected according to (6).
- 5: Using these subsets and (4), populate the matrix with values:

$$(\tilde{K})_{ij} = \tilde{k}_{\text{set}}(\mathbf{X}^{(i)}, \mathbf{X}^{(j)}; \mathbf{w}, \pi, L).$$

6: **return** a computed matrix  $\tilde{K}$

---

To impose an ordering of the elements, we use a *random scalar projection*  $\mathbf{w} \in \mathbb{R}^d$  such that the elements of  $\mathbf{w}$  are drawn from the standard normal distribution. If the scalar projections of each  $\mathbf{x}_i$  are computed, this produces the set of scalar values  $\{\mathbf{w}^\top \mathbf{x}_1, \dots, \mathbf{w}^\top \mathbf{x}_m\}$ , which can be sorted to generate an ordered list of

$$[\ell_1, \dots, \ell_m], \quad \mathbf{w}^\top \mathbf{x}_{\ell_1} \leq \dots \leq \mathbf{w}^\top \mathbf{x}_{\ell_m}, \tag{5}$$

for an ordering of distinct indices  $\ell_1, \dots, \ell_m \in [m]$ . Ties between  $\mathbf{w}^\top \mathbf{x}_i$  values can be dealt with arbitrarily. The function  $\pi$  then is simply a random bijection of the integers  $[m]$  onto themselves. Using this, we can sample  $L$  vectors from  $\mathbf{X}^{(i)}$ :

$$\tilde{\mathbf{X}}^{(i)} = \{\mathbf{x}_{\ell_j} \mid \ell_j = \pi(j) \text{ for } j \in [L]\}. \tag{6}$$

This process, given  $\mathbf{w}$ ,  $\pi$ , and  $L$ , is sufficient for computing  $k_{\text{set}}$ , as presented in Algorithm 1.

**3.2 Properties of the approximation**

The covariance matrix for this approximation of the set kernel, which we denote by  $(\tilde{K})_{ij} = \tilde{k}_{\text{set}}(\mathbf{X}^{(i)}, \mathbf{X}^{(j)}; \mathbf{w}, \pi, L)$ , should approximate the full version of covariance matrix,  $K$  from (3). Because of the random structure introduced in Sect. 3.1, the matrix  $\tilde{K}$  will be random. This will be addressed in Theorem 1, but for now,  $\tilde{K}$  represents a single realization of that random variable, not the random variable itself. To be viable, this approximation must satisfy the following requirements:

**Property 1** *The approximation satisfies pairwise symmetry:*

$$\tilde{k}_{\text{set}}(\mathbf{X}^{(i)}, \mathbf{X}^{(j)}; \mathbf{w}, \pi, L) = \tilde{k}_{\text{set}}(\mathbf{X}^{(j)}, \mathbf{X}^{(i)}; \mathbf{w}, \pi, L). \tag{7}$$

Since  $\tilde{\mathbf{X}}^{(i)}$  is uniquely defined given  $\mathbf{w}, \pi, L$ , this simplifies to  $k_{\text{set}}(\tilde{\mathbf{X}}^{(i)}, \tilde{\mathbf{X}}^{(j)}) = k_{\text{set}}(\tilde{\mathbf{X}}^{(j)}, \tilde{\mathbf{X}}^{(i)})$ , which is true because  $k_{\text{set}}$  is symmetric.

**Property 2** *The “ordering” of the elements in the sets  $\mathbf{X}^{(i)}, \mathbf{X}^{(j)}$  should not matter when computing  $\tilde{k}_{\text{set}}$ . Indeed, because (5) enforces ordering based on  $\mathbf{w}$ , and not*

whatever arbitrary indexing is imposed in defining the elements of the set, the kernel will be permutation-invariant.

**Property 3** The kernel approximation (4) reduces to computing  $k_{\text{set}}$  on a lower cardinality version of the data (with  $L$  elements selected from  $m$ ). Because  $k_{\text{set}}$  is positive-definite on these  $L$ -element sets, we know that  $\bar{k}_{\text{set}}$  is also positive-definite.

**Property 4** Since the approximation method aims to choose subsets of input sets, the computational cost becomes lower than the original formulation.

Missing from these four properties is a statement regarding the quality of the approximation. We address this in Theorems 1 and 2, though we first start by stating Lemma 2.

**Lemma 2** Suppose there are two sets  $\mathbf{X}, \mathbf{Y} \in \mathcal{X}_{\text{set}}$ . Without loss of generality, let  $\mathbf{X}^{(i)}$  and  $\mathbf{Y}^{(j)}$  denote the  $i$ th and  $j$ th of  $\binom{m}{L}$  possible subsets containing  $L$  elements of  $\mathbf{X}$  and  $\mathbf{Y}$ , respectively, in an arbitrary ordering. For  $L \in [m]$ ,

$$\sum_{i=1}^{\binom{m}{L}} \sum_{j=1}^{\binom{m}{L}} \sum_{a=1}^L \sum_{b=1}^L k(\bar{\mathbf{x}}_a^{(i)}, \bar{\mathbf{y}}_b^{(j)}) = \frac{L^2 \binom{m}{L}^2}{m^2} \sum_{c=1}^m \sum_{d=1}^m k(\mathbf{x}_c, \mathbf{y}_d), \tag{8}$$

where  $\bar{\mathbf{x}}_a^{(i)}$  and  $\bar{\mathbf{y}}_b^{(j)}$  are the  $a$ th and  $b$ th elements of  $\mathbf{X}^{(i)}$  and  $\mathbf{Y}^{(j)}$ , respectively, in an arbitrary ordering.

**Proof** We can rewrite the original summation in a slightly more convoluted form, as

$$\begin{aligned} & \sum_{i=1}^{\binom{m}{L}} \sum_{j=1}^{\binom{m}{L}} \sum_{a=1}^L \sum_{b=1}^L k(\bar{\mathbf{x}}_a^{(i)}, \bar{\mathbf{y}}_b^{(j)}) \\ &= \sum_{i=1}^{\binom{m}{L}} \sum_{j=1}^{\binom{m}{L}} \sum_{a=1}^L \sum_{b=1}^L \sum_{c=1}^m \sum_{d=1}^m k(\bar{\mathbf{x}}_a^{(i)}, \bar{\mathbf{y}}_b^{(j)}) I_{\bar{\mathbf{x}}_a^{(i)}, \bar{\mathbf{y}}_b^{(j)}}(\mathbf{x}_c, \mathbf{y}_d) \\ &= \sum_{c=1}^m \sum_{d=1}^m \left[ \sum_{i=1}^{\binom{m}{L}} \sum_{j=1}^{\binom{m}{L}} \sum_{a=1}^L \sum_{b=1}^L k(\bar{\mathbf{x}}_a^{(i)}, \bar{\mathbf{y}}_b^{(j)}) I_{\bar{\mathbf{x}}_a^{(i)}, \bar{\mathbf{y}}_b^{(j)}}(\mathbf{x}_c, \mathbf{y}_d) \right], \end{aligned} \tag{9}$$

where  $I_{\bar{\mathbf{x}}_a^{(i)}, \bar{\mathbf{y}}_b^{(j)}}(\mathbf{x}_c, \mathbf{y}_d) = 1$  if  $\bar{\mathbf{x}}_a^{(i)} = \mathbf{x}_c$  and  $\bar{\mathbf{y}}_b^{(j)} = \mathbf{y}_d$ , and 0 otherwise. As these are finite summations, they can be safely reordered.

The symmetry in the structure and evaluation of the summation implies that as each  $\mathbf{x}_c$  quantity will be paired with each  $\mathbf{y}_d$  quantity the same number of times. Therefore, we need only consider the number of times that these quantities appear.

We recognize that this summation follows a pattern related to Pascal’s triangle. Among the  $\binom{m}{L}$  possible subsets  $\bar{\mathbf{x}}$  of  $\mathbf{X}$ , only the fraction  $L/m$  of those contain the quantity  $\mathbf{x}_c$  for all  $c \in [m]$  (irrespective of how that entry may be denoted in  $\bar{\mathbf{x}}_a^{(i)}$  terminology). Because of the symmetry mentioned above, each of those  $\mathbf{x}_c$  quantities is paired with each of the  $\mathbf{y}_d$  quantities the same  $\frac{L}{m} \binom{m}{L}$  number of times. This result implies that

$$\sum_{i=1}^m \sum_{j=1}^m \sum_{a=1}^L \sum_{b=1}^L k(\bar{\mathbf{x}}_a^{(i)}, \bar{\mathbf{y}}_b^{(j)}) I_{\bar{\mathbf{x}}_a^{(i)}, \bar{\mathbf{y}}_b^{(j)}}(\mathbf{x}_c, \mathbf{y}_d) = \frac{L^2 \binom{m}{L}^2}{m^2} k(\mathbf{x}_c, \mathbf{y}_d), \tag{10}$$

where  $I_{\bar{\mathbf{x}}_a^{(i)}, \bar{\mathbf{y}}_b^{(j)}}(\mathbf{x}_c, \mathbf{y}_d) = 1$  if  $\bar{\mathbf{x}}_a^{(i)} = \mathbf{x}_c$  and  $\bar{\mathbf{y}}_b^{(j)} = \mathbf{y}_d$ , and 0 otherwise. Substituting (10) into the bracketed quantity in (9) above completes the proof.  $\square$

We start by introducing the notation  $W$  and  $\Pi$  to be random variables such that  $W \sim \mathcal{N}(0, \mathbf{I}_d)$  and  $\Pi$  is a uniformly random permutation of the integers between 1 and  $m$ . These are the distributions defining the  $\mathbf{w}$  and  $\pi$  quantities described above. With this, we note that  $\tilde{k}_{\text{set}}(\mathbf{X}, \mathbf{Y}; W, \Pi, L)$  is a random variable.

We also introduce the notation  $\sigma_L(\mathbf{X})$  to be the distribution of random subsets of  $\mathbf{X}$  with  $L$  elements selected without replacement, the outcome of the subset selection from Sect. 3.1. This notation allows us to write the quantities

$$\mathbb{E}_{W, \Pi}[\tilde{k}_{\text{set}}(\mathbf{X}, \mathbf{Y}; W, \Pi, L)] = \mathbb{E}_{\bar{\mathbf{X}}, \bar{\mathbf{Y}}}[k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}})] \equiv \mathbb{E}[k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}})], \tag{11}$$

$$\text{Var}_{W, \Pi}[\tilde{k}_{\text{set}}(\mathbf{X}, \mathbf{Y}; W, \Pi, L)] = \text{Var}_{\bar{\mathbf{X}}, \bar{\mathbf{Y}}}[k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}})] \equiv \text{Var}[k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}})], \tag{12}$$

for  $\bar{\mathbf{X}} \sim \sigma_L(\mathbf{X})$ ,  $\bar{\mathbf{Y}} \sim \sigma_L(\mathbf{Y})$ . We have dropped the random variables from the expectation and variance definitions for ease of notation.

**Theorem 1** *Suppose that we are given two sets  $\mathbf{X}, \mathbf{Y} \in \mathcal{X}_{\text{set}}$  and  $L \in \mathbb{Z}_+$ . Suppose, furthermore, that  $\mathbf{w}$  and  $\pi$  can be generated randomly as defined in Sect. 3.1 to form subsets  $\tilde{\mathbf{X}}$  and  $\tilde{\mathbf{Y}}$ . The value of  $\tilde{k}_{\text{set}}(\mathbf{X}, \mathbf{Y}; \mathbf{w}, \pi, L)$  is an unbiased estimator of the value of  $k_{\text{set}}(\mathbf{X}, \mathbf{Y})$ .*

**Proof** Our goal is to show that  $\mathbb{E}[k_{\text{set}}(\tilde{\mathbf{X}}, \tilde{\mathbf{Y}})] = k_{\text{set}}(\mathbf{X}, \mathbf{Y})$ , where  $\mathbb{E}[k_{\text{set}}(\tilde{\mathbf{X}}, \tilde{\mathbf{Y}})]$  is defined in (11).

We first introduce an extreme case:  $L = m$ . If  $L = m$ , the subsets we are constructing are the full sets, i.e.,  $\sigma_m(\mathbf{X})$  contains only one element,  $\mathbf{X}$ . Thus,  $\tilde{k}_{\text{set}}(\mathbf{X}, \mathbf{Y}; W, \Pi, m) = k_{\text{set}}(\mathbf{X}, \mathbf{Y})$  is not a random variable.

For  $1 \leq L < m$ , we compute this expected value from the definition (with some abuse of notation):

$$\mathbb{E}[k_{\text{set}}(\tilde{\mathbf{X}}, \tilde{\mathbf{Y}})] = \sum_{\mathbf{X}, \mathbf{Y}} k_{\text{set}}(\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}) p(\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}). \tag{13}$$



There are  $\binom{m}{L}$  subsets, all of which could be indexed (arbitrarily) as  $\bar{\mathbf{X}}^{(i)}$  for  $1 \leq i \leq \binom{m}{L}$ . The probability mass function is uniform across all subsets, meaning that  $p(\bar{\mathbf{X}} = \bar{\mathbf{X}}^{(i)}, \bar{\mathbf{Y}} = \bar{\mathbf{Y}}^{(i)}) = 1/\binom{m}{L}^2$ . Using this, we know

$$\mathbb{E}[k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}})] = \sum_{i=1}^{\binom{m}{L}} \sum_{j=1}^{\binom{m}{L}} k_{\text{set}}(\bar{\mathbf{X}}^{(i)}, \bar{\mathbf{Y}}^{(j)}) \frac{1}{\binom{m}{L}^2}. \tag{14}$$

We apply (2) to see that

$$k_{\text{set}}(\bar{\mathbf{X}}^{(i)}, \bar{\mathbf{Y}}^{(j)}) = \frac{1}{L^2} \sum_{a=1}^L \sum_{b=1}^L k(\bar{\mathbf{x}}_a^{(i)}, \bar{\mathbf{y}}_b^{(j)}), \tag{15}$$

following the notational conventions used above. The expectation involves four nested summations,

$$\mathbb{E}[k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}})] = \frac{1}{L^2 \binom{m}{L}^2} \sum_{i=1}^{\binom{m}{L}} \sum_{j=1}^{\binom{m}{L}} \sum_{a=1}^L \sum_{b=1}^L k(\bar{\mathbf{x}}_a^{(i)}, \bar{\mathbf{y}}_b^{(j)}). \tag{16}$$

We utilize Lemma 2 to rewrite this as

$$\mathbb{E}[k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}})] = \frac{1}{L^2 \binom{m}{L}^2} \frac{L^2 \binom{m}{L}^2}{m^2} \sum_{c=1}^m \sum_{d=1}^m k(\mathbf{x}_c, \mathbf{y}_d) = \frac{1}{m^2} \sum_{c=1}^m \sum_{d=1}^m k(\mathbf{x}_c, \mathbf{y}_d). \tag{17}$$

□

**Theorem 2** Under the same conditions as in Theorem 1, suppose that  $k(\mathbf{x}, \mathbf{x}') \geq 0$  for all  $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ . The variance of  $\tilde{k}_{\text{set}}(\mathbf{X}, \mathbf{Y}; \mathbf{w}, \pi, L)$  is bounded by a function of  $m, L$  and  $k_{\text{set}}(\mathbf{X}, \mathbf{Y})$ :

$$\text{Var}[\tilde{k}_{\text{set}}(\mathbf{X}, \mathbf{Y}; \mathbf{w}, \pi, L)] \leq \left(\frac{m^4}{L^4} - 1\right) k_{\text{set}}(\mathbf{X}, \mathbf{Y})^2. \tag{18}$$

**Proof** The variance of  $k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}})$ , defined in (12), is computed as

$$\begin{aligned} \text{Var}[k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}})] &= \mathbb{E}\left[\left(k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}}) - \mathbb{E}[k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}})]\right)^2\right] \\ &= \mathbb{E}[k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}})^2] + k_{\text{set}}(\mathbf{X}, \mathbf{Y})^2 - 2k_{\text{set}}(\mathbf{X}, \mathbf{Y})\mathbb{E}[k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}})] \\ &= \mathbb{E}[k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}})^2] - k_{\text{set}}(\mathbf{X}, \mathbf{Y})^2, \end{aligned} \tag{19}$$

where Theorem 1 is invoked to produce the final line. Using (14) and (15), we can express the first term of (19) as

$$\mathbb{E}[k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}})^2] = \sum_{i=1}^m \sum_{j=1}^m \left( \frac{1}{L^2} \sum_{a=1}^L \sum_{b=1}^L k(\bar{\mathbf{x}}_a^{(i)}, \bar{\mathbf{x}}_b^{(j)}) \right)^2 \frac{1}{\binom{m}{L}^2}. \tag{20}$$

At this point, we invoke the fact that  $k(\mathbf{x}, \mathbf{x}') \geq 0$  to state

$$0 \leq \sum_{a=1}^L \sum_{b=1}^L k(\bar{\mathbf{x}}_a^{(i)}, \bar{\mathbf{x}}_b^{(j)}) \leq \sum_{a=1}^m \sum_{b=1}^m k(\mathbf{x}_a, \mathbf{x}_b), \tag{21}$$

which is true because the summation to  $m$  terms contains all of the elements in the summation to  $L$  terms, as well as other (nonnegative) elements. Using this, we can bound (20) by

$$\begin{aligned} \mathbb{E}[k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}})^2] &\leq \sum_{i=1}^m \sum_{j=1}^m \left( \frac{1}{L^2} \sum_{a=1}^m \sum_{b=1}^m k(\mathbf{x}_a, \mathbf{x}_b) \right)^2 \frac{1}{\binom{m}{L}^2} \\ &= \frac{m^4}{\binom{m}{L}^2 L^4} \sum_{i=1}^m \sum_{j=1}^m \left( \frac{1}{m^2} \sum_{a=1}^m \sum_{b=1}^m k(\mathbf{x}_a, \mathbf{x}_b) \right)^2 \\ &= \frac{m^4}{\binom{m}{L}^2 L^4} \left( \binom{m}{L} \right)^2 k_{\text{set}}(\mathbf{X}, \mathbf{Y})^2 = \frac{m^4}{L^4} k_{\text{set}}(\mathbf{X}, \mathbf{Y})^2. \end{aligned} \tag{22}$$

Therefore, with (22), (19) can be written as

$$\text{Var}[k_{\text{set}}(\bar{\mathbf{X}}, \bar{\mathbf{Y}})] \leq \frac{m^4}{L^4} k_{\text{set}}(\mathbf{X}, \mathbf{Y})^2 - k_{\text{set}}(\mathbf{X}, \mathbf{Y})^2 = \left( \frac{m^4}{L^4} - 1 \right) k_{\text{set}}(\mathbf{X}, \mathbf{Y})^2, \tag{23}$$

which concludes this proof. □

The restriction  $k(\mathbf{x}, \mathbf{x}') \geq 0$  is satisfied by many standard covariance kernels (such as the Gaussian, the Matérn family and the multiquadric) as well as some more interesting choices (such as the Wendland or Wu families of compactly supported kernels). It does, however, exclude some oscillatory kernels such as the Poisson kernel as well as kernels defined implicitly which may have an oscillatory behavior. More discussion on different types of kernels and their properties can be found in the kernel literature (Fasshauer and McCourt 2015).

By Theorem 2, we can naturally infer the fact that the upper bound of the variance decreases quickly, as  $L$  is close to  $m$ .

### 3.3 Bayesian optimization over sets

For Bayesian optimization over  $\mathcal{X}_{\text{set}}$ , the goal is to identify the set  $\mathbf{X} \in \mathcal{X}_{\text{set}}$  such that a given function  $f : \mathcal{X}_{\text{set}} \rightarrow \mathbb{R}$  is minimized. As shown in Algorithm 2, Bayesian optimization over sets follows similar steps as laid out in Sect. 2, except that it involves the space of set inputs and requires a surrogate function on  $\mathcal{X}_{\text{set}}$ . As we have already indicated, we plan to use a GP surrogate function, with prior covariance defined either with (2) or (4) and a Matérn 5/2 inner kernel  $k$ .

A GP model requires computation on the order of  $\mathcal{O}(n^3)$  at the  $n$ th step of the Bayesian optimization because the K matrix must be inverted. Compared to the complexity for computing a full version of the set kernel  $\mathcal{O}(n^2m^2d)$ , the complexity of computing the inverse is smaller if roughly  $m \geq n$  (that is, computing the matrix can be as costly or more costly than inverting it). Because Bayesian optimization is efficient sampling-based global optimization,  $n$  is small and the situation  $m \geq n$  is reasonable. Therefore, the computation reduction by our approximation can be effective in reducing complexity of all steps for Bayesian optimization over sets.

In addition, since the Cholesky decomposition, instead of matrix inverse is widely used to compute posterior mean and variance functions (Rasmussen and Williams 2006), the time complexity for inverting a covariance matrix can be reduced. But still, if  $m$  is relatively large, our approximation is effective. In this paper, we compute GP surrogate using the Cholesky decomposition. See the effects of  $L$  in Sect. 4.1.

---

#### Algorithm 2 Bayesian Optimization over Sets

---

**Input:** A domain  $\mathcal{X}_{\text{set}}$ , a function  $f : \mathcal{X}_{\text{set}} \rightarrow \mathbb{R}$ , a budget  $T \in \mathbb{Z}_+$ .

**Output:** Best acquired set  $\mathbf{X}^\dagger$ .

- 1: Choose an initial point  $\mathbf{X}^{(1)}$  randomly from  $\mathcal{X}_{\text{set}}$  and evaluate  $y_1 = f(\mathbf{X}^{(1)}) + \epsilon_1$ .
  - 2: **for**  $k$  from 1 to  $T - 1$  **do**
  - 3:   Fit the surrogate model  $s_k$  to all data  $\{(\mathbf{X}^{(i)}, y_i)\}_{i=1}^k$ .
  - 4:   Compute the acquisition function  $a_k$  from  $s_k$ .
  - 5:   Identify  $\mathbf{X}^{(k+1)} = \arg \max_{\mathbf{X} \in \mathcal{X}_{\text{set}}} a_k(\mathbf{X})$ .
  - 6:   Evaluate  $y_{k+1} = f(\mathbf{X}^{(k+1)}) + \epsilon_{k+1}$ .
  - 7: **end for**
  - 8: **return**  $\mathbf{X}^\dagger = \mathbf{X}^{(i)}$  if  $y_i = \max_{j \in [T]} y_j$
- 

### 3.4 Acquisition function optimization over sets

An acquisition function optimization step is one of the primary steps in Bayesian optimization, because this step is for finding an acquired example, which exhibits the highest potential of the global optimum. Compared to generic vector-input Bayesian optimization methods, our Bayesian optimization over sets needs to find a query set  $\mathbf{X}$  on  $\mathcal{X}_{\text{set}}$ , optimizing the acquisition function over sets.

Since  $m$  is fixed when we find  $\mathbf{X}$ , off-the-shelf optimization methods such as L-BFGS-B (Liu and Nocedal 1989), DIRECT (Jones et al. 1993), and CMA-ES (Hansen 2016) can be employed where a set is treated as a concatenated vector. However, because of the symmetry of the space  $\mathcal{X}_{\text{set}}$ , these common optimization methods search more of the space than is required. For instance, if we optimize the function on  $\{1, 2, 3\}$ , we would not need to consider such the sets  $\{1, 3, 2\}$ ,  $\{2, 1, 3\}$ ,  $\dots$ ,  $\{3, 2, 1\}$ . Therefore, we suggest a constrained acquisition

function optimization with rejection sampling algorithm. The sample complexity of this method decreases by a factor of  $m!$ .

First of all, we sample all instances of the sets to initialize acquisition function optimization from uniform distribution. By the rejection sampling, some of the sets sampled are rejected if each of them is located outside of the symmetric region. The optimization method (i.e., CMA-ES) finds the optimal set of the acquisition function, starting from those initial sets selected from the aforementioned step. In the optimization step, our optimization strategy forces every optimization result to locate in the symmetric search space. Finally, we pick the best one among the converged sets as the next set to observe.

### 3.5 Regret analysis

To produce a regret bound on Bayesian optimization with our approximation as well as one of Bayesian optimization with the set kernel, we follow the framework to prove the regret bounds of the multi-armed bandit-like acquisition function in Srinivas et al. (2010). To simplify the analysis, we discuss only the Matérn kernel for  $\nu > 2$  as a base kernel for GP.

Inspired by Kandasamy et al. (2015), we assume that our objective function  $f$  can be expressed as a summation of functions over instances and they can be collected to a function that can take a single vector with the i.i.d. property of instances in a set:

$$f(\mathbf{X}) = \frac{1}{m} \sum_{i=1}^m g(\mathbf{x}_i) = h([\mathbf{x}_1, \dots, \mathbf{x}_m]), \tag{24}$$

which implies that  $h$  can take  $md$ -dimensional concatenated inputs. Thus, we can state our Bayesian optimization with the set kernel follows the cumulative regret bound proposed by Srinivas et al. (2010) and Kandasamy et al. (2015). Given  $T$  available sets  $\{\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(T)}\}$ , let  $\delta \in (0, 1)$  and trade-off hyperparameter for GP-UCB  $\beta_t \in \mathcal{O}(md \log t)$ . A cumulative regret bound  $R_T$  is

$$R_T = \mathcal{O}(md2^d T^{\frac{\nu+d(d+1)}{2\nu+d(d+1)}}), \tag{25}$$

with the probability at least  $1 - \delta$ , under the mild assumptions: (i)  $k_{\text{set}}(\mathbf{X}, \mathbf{X}') \leq 1$ ; (ii) the bounded reproducing kernel Hilbert space (RKHS) norm  $\|f\|_{k_{\text{set}}} < b$  where  $b > 0$ . Moreover,  $\tilde{k}_{\text{set}}(\mathbf{X}, \mathbf{X}'; \mathbf{w}, \pi, L) \leq 1$  and  $\|f\|_{\tilde{k}_{\text{set}}} < b$  are naturally satisfied by their definitions.

By Theorems 1 and 2, we can define Corollary 1, which is related to the aforementioned regret bound.

**Corollary 1** *Suppose that  $0 < \delta \ll 1$ .  $R_T^{(L)}$  is a cumulative regret computed by the approximate set kernel with  $L$  of  $m$  instances in each set. By Theorems 1 and 2,  $\mathbb{E}_L[R_T^{(L)}] = R_T$ , is satisfied with the probability at least  $1 - \delta$ .*

**Proof** Since our approximation is an unbiased estimator of original set kernel  $k_{\text{set}}$  with bounded variance as discussed in Theorems 1 and 2, the expectation of cumulative regrets with the approximation is equal to  $R_T$ . □

## 4 Experiments

In this section, we present various experimental results, to show unique applications of our method as well as the motivating problems. First, we conduct our method on the experiments regarding set kernel approximation and constrained acquisition function optimization, in order to represent the effectiveness of our proposed method. Then, we optimize two synthetic functions and clustering algorithm initialization, which take a set as an input. Finally, we present the experimental results on active nearest neighbor search for point clouds.

We define the application-agnostic baseline methods, Vector and Split:

**Vector** A standard Bayesian optimization is performed over a  $md$ -dimensional space where, at the  $n$ th step, the available data  $\mathfrak{X}_n \in \mathcal{X}_{\text{set}}^n$  is vectorized to  $[\mathbf{x}_1, \dots, \mathbf{x}_n]$  for  $\mathbf{x}_i \in \mathbb{R}^{md}$  with associated function values. At each step, the vectorized next location  $\mathbf{x}_{n+1}$  is converted into a set  $\mathbf{X}_{n+1}$ .

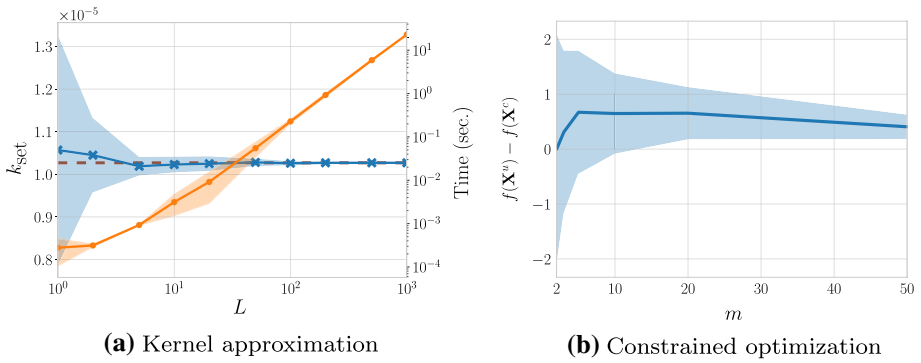
**Split** Individual Bayesian optimization strategies are executed on the  $m$  components comprising  $\mathcal{X}$ . At the  $n$ th step, the available data  $\mathfrak{X}_n \in \mathcal{X}_{\text{set}}^n$  is decomposed into  $m$  sets of data, the  $i$ th of which consists of  $[\mathbf{x}_1^{(i)}, \dots, \mathbf{x}_n^{(i)}]$  with associated data. The  $m$  vectors produced during each step of the optimization are then collected to form  $\mathbf{X}_{n+1}$  at which to evaluate  $f$ .

For Vector and Split baselines, to satisfy the permutation-invariance property, we determine the order of elements in a set as the ascending order by the  $l^2$  norm of the elements.

We use Gaussian process regression (Rasmussen and Williams 2006) with a set kernel or Matérn 5/2 kernel as a surrogate function. Because computing the inverse of covariance matrix needs heavy computations, we employ the Cholesky decomposition instead (Rasmussen and Williams 2006). For the experiments with set kernel, Matérn 5/2 kernel is used as a base kernel. All Gaussian process regression models are optimized by marginal likelihood maximization with the BFGS optimizer, to find kernel hyperparameters. We use Gaussian process upper confidence bound criterion (Srinivas et al. 2010) as an acquisition function for all experiments. CMA-ES (Hansen 2016) and its constrained version are applied in optimizing the acquisition function. Furthermore, five initial points are given to start a single round of Bayesian optimization. Unless otherwise specified, most of experiments are repeated 10 times. For the results on execution time, all the results include the time consumed in evaluating a true function. All results are run via CPU computations. All implementations which will be released as open source project are written in Python. Thanks to Pedregosa et al. (2011), we use `scikit-learn` in many parts of our implementations.

### 4.1 Set kernel approximation

We study the effect of  $L$  for the set kernels. Using a set generated from the standard normal distribution, which has 1000 50-dimensional instances, we observe the effects of  $L$  as shown in Fig. 2a.  $k_{\text{set}}$  converges to the true value as  $L$  increases, and the variance of  $k_{\text{set}}$  value is large when  $L$  is small, as discussed in Sect. 3.2. Moreover, the consumed time increases as  $L$  increases. We use Matérn 5/2 kernel as a base kernel. Table 1 shows the effects of  $L$  for set kernels. As  $L$  increases,  $k_{\text{set}}$  value is converged to the true value and execution time increases.



**Fig. 2** Results on the effects of  $L$  for set kernels and constrained optimization for acquisition functions. The mean and standard deviation of each quantity are plotted, computed over 10 trials. (left) The lines with  $\times$  and  $\circ$  indicate kernel values and consumed times, respectively. The dashed line is the true set kernel value. (right)  $X^u$  and  $X^c$  are the acquired sets obtained by the acquisition function optimization w/o and w/constraints, respectively

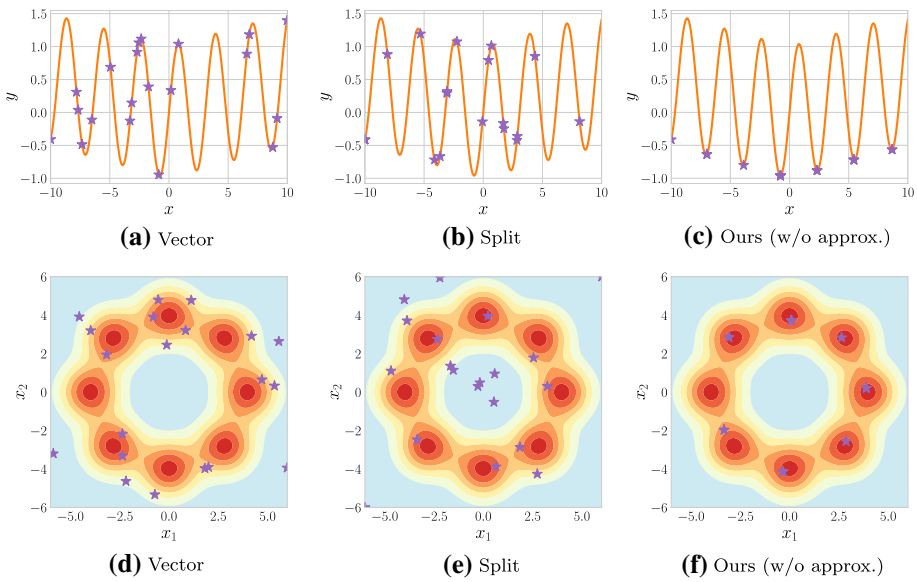
**Table 1** The effects of  $L$  for set kernels

$L$	$k_{\text{set}}$	Time (s)
1	$1.057 \times 10^{-5} \pm 1.387 \times 10^{-6}$	$(2.776 \pm 0.877) \times 10^{-4}$
2	$1.045 \times 10^{-5} \pm 4.512 \times 10^{-7}$	$(3.131 \pm 0.118) \times 10^{-4}$
5	$1.019 \times 10^{-5} \pm 1.168 \times 10^{-7}$	$(9.283 \pm 0.089) \times 10^{-4}$
10	$1.023 \times 10^{-5} \pm 9.980 \times 10^{-8}$	$(3.133 \pm 0.822) \times 10^{-3}$
20	$1.025 \times 10^{-5} \pm 9.305 \times 10^{-8}$	$(9.150 \pm 3.169) \times 10^{-3}$
50	$1.028 \times 10^{-5} \pm 4.955 \times 10^{-8}$	$(5.487 \pm 1.075) \times 10^{-2}$
100	$1.026 \times 10^{-5} \pm 3.030 \times 10^{-8}$	$(2.282 \pm 0.134) \times 10^{-1}$
200	$1.027 \times 10^{-5} \pm 1.961 \times 10^{-8}$	$(9.253 \pm 0.490) \times 10^{-1}$
500	$1.027 \times 10^{-5} \pm 1.372 \times 10^{-8}$	$(5.887 \pm 0.159) \times 10^0$
1000(= $m$ )	$1.027 \times 10^{-5}$	$(2.295 \pm 0.027) \times 10^1$

All settings follow the settings in Fig. 2a. The numerical results are rounded to the three decimals, to show the effects precisely

### 4.2 Constrained acquisition function optimization

We demonstrate the effects of the constrained acquisition function optimization, compared to the vanilla optimization method that concatenates a set to a single vector. In this paper, we use CMA-ES (Hansen 2016) as an acquisition function optimization method, which is widely used in Bayesian optimization (Bergstra et al. 2011; Wang et al. 2014). As we mentioned in Sect. 3.4, the constrained CMA-ES is more sample-efficient than the vanilla CMA-ES. Fig. 2b (i.e., a minimization problem) represents that the function values determined by the constrained method are always smaller than the values by the unconstrained method, because we fix the number of initial samples (i.e., 5). Moreover, the variance of them decreases, as  $m$  is large. For this experiment, we measure the acquisition performance where 20 fixed historical observations are given, and Synthetic 1 described in Sect. 4.3 is used. Note that the kernel approximation is not applied.



**Fig. 3** Examples of one of the best acquisition results (i.e., purple stars indicate instances in the acquired set) via Vector, Split, and Ours (w/o approximation). For Synthetic 1 (first row) and Synthetic 2 (second row),  $m$  is set to 20 (Color figure online)

### 4.3 Synthetic functions

We test two synthetic functions to show Bayesian optimization over sets is a valid approach to find an optimal set that minimizes an objective function  $f : \mathcal{X}_{\text{set}} \rightarrow \mathbb{R}$ . In each setting, there is an auxiliary function  $g : \mathcal{X} \rightarrow \mathbb{R}$ , and  $f$  is defined as  $f(\mathbf{X}) = \frac{1}{m} \sum_{i=1}^m g(\mathbf{x}_i)$ . The  $g$  functions are designed to be multi-modal, giving the opportunity for the set  $\mathbf{X}$  to contain  $\mathbf{x}_i$  values from each of the modes in the domain. Additionally, as is expected,  $f$  is permutation invariant (any ordering imposed on the elements of  $\mathbf{x}$  is immaterial).

**Synthetic 1** We consider  $d = 1, m = 20$  and choose  $g$  to be a simple periodic function:

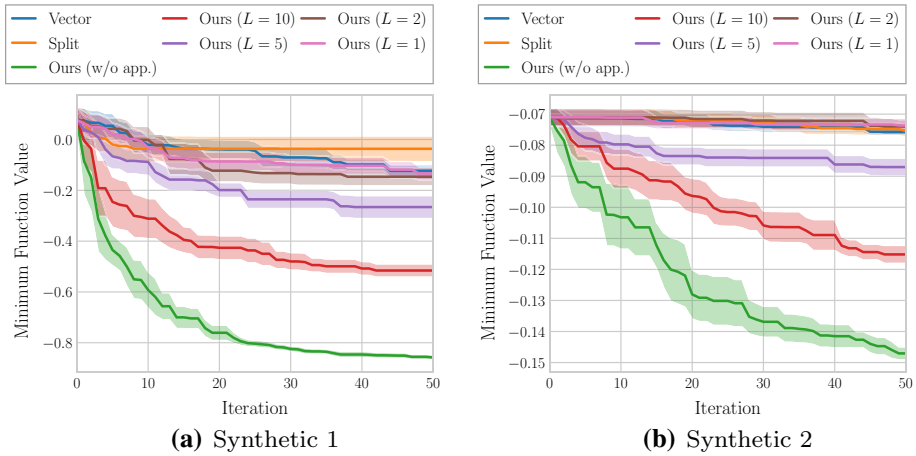
$$g(\mathbf{x}) = \sin(2\|\mathbf{x}\|_2) + |0.05\|\mathbf{x}\|_2|. \tag{26}$$

**Synthetic 2** We consider  $d = 2, m = 20$  and a  $g$  function which is the sum of probability density functions:

$$g(\mathbf{x}) = - \sum_{i=1}^8 p(\mathbf{x}; \mu_i, \Sigma_i), \tag{27}$$

where  $p$  is the normal density function with  $\mu_i$  depicted in Fig. 3 and  $\Sigma_i = \mathbf{I}_2$ .

As shown in Fig. 3, both of these functions have a clear multimodal structure, allowing for optimal sets to contain points which are clustered in a single local minima or to be spread out through the domain in several local minima. Fig. 4 shows that Vector and Split strategies have difficulty optimizing the functions. On the other hand, our proposed method



**Fig. 4** Results on optimizing two synthetic functions. As presented in Fig. 3,  $m$  is set to 20. All experiments are repeated 10 times

**Table 2** Convergence quality and its execution time on two synthetic functions where  $m = 20$

$L$	Synthetic 1		Synthetic 2	
	Minimum	Time ( $10^3$ s)	Minimum	Time ( $10^3$ s)
1	$-0.133 \pm 0.052$	$5.682 \pm 0.279$	$-0.074 \pm 0.004$	$6.468 \pm 0.130$
2	$-0.146 \pm 0.071$	$8.003 \pm 0.156$	$-0.074 \pm 0.005$	$9.768 \pm 0.136$
5	$-0.266 \pm 0.090$	$16.299 \pm 0.253$	$-0.087 \pm 0.006$	$20.451 \pm 0.033$
10	$-0.516 \pm 0.052$	$31.010 \pm 0.810$	$-0.115 \pm 0.006$	$40.099 \pm 0.085$
20(= $m$ )	$-0.858 \pm 0.019$	$91.589 \pm 2.339$	$-0.147 \pm 0.004$	$85.266 \pm 0.297$

All settings follow the settings in Figs. 3 and 4

finds optimal outcomes more effectively.<sup>3</sup> We study the impact of  $L$  when optimizing these two synthetic functions; a smaller  $L$  should yield faster computations, but also a worse approximation  $\tilde{K}$  to the true  $K$  matrix (when  $L = m$ ).

Table 2 represents a convergence quality and its execution time for the synthetic functions defined in this work. As expected, the execution time decreases as  $L$  decreases.

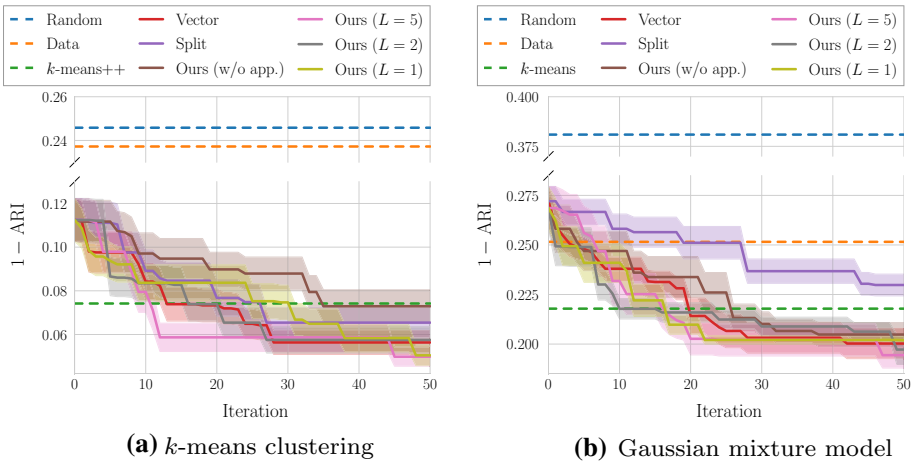
#### 4.4 Clustering algorithm initialization

We initialize clustering algorithms for dataset  $\mathcal{P} = [\mathbf{p}_1, \dots, \mathbf{p}_N]$  with Bayesian optimization over sets. For these experiments, we add four additional baselines for clustering algorithms:

*Random* This baseline randomly draws  $k$  points from a compact space  $\subset \mathbb{R}^d$ .

<sup>3</sup> While not our concern here, it is possible that some amount of distance between points in the set  $\mathbf{X}$  would be desired. If that were the case, such a desire could be enforced in the function  $f$ .





**Fig. 5** Results on initializing clustering algorithms: *k*-means clustering and Gaussian mixture model for synthetic datasets

*Data* This baseline randomly samples  $k$  points from a dataset  $\mathcal{P}$ . It is widely used in initializing a clustering algorithm.

*(k-means only)* *k*-means++ (Arthur and Vassilvitskii 2007) This is a method for *k*-means clustering with the intuition that spreading out initial cluster centers is better than the Data baseline.

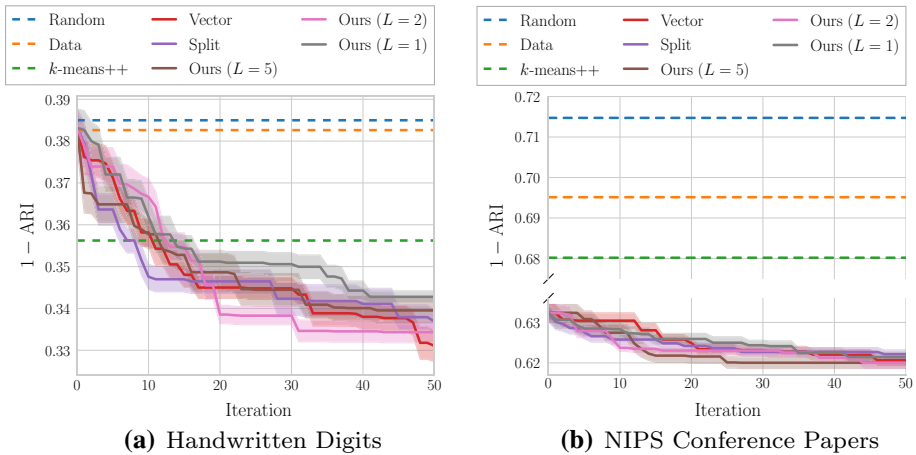
*(GMM only)* *k*-means This baseline sets initial cluster centers as the results of *k*-means clustering.

To fairly compare the baselines to our methods, the baselines are trained by the whole datasets without splitting. To compare with the baselines fairly, Random, Data, *k*-means++ (Arthur and Vassilvitskii 2007), *k*-means are run 1000 times. In Bayesian optimization settings, we split a dataset to training (70%) and test (30%) datasets. After finding the converged cluster centers  $\{\mathbf{c}_1, \dots, \mathbf{c}_k\}$  with training dataset, the adjusted Rand index (ARI) is computed by test dataset. The algorithms are optimized over  $1 - \text{ARI}$ . All clustering models are implemented using `scikit-learn` (Pedregosa et al. 2011).

We test two clustering algorithms for synthetic datasets: (i) *k*-means clustering and (ii) Gaussian mixture model (GMM). In addition, two real-world datasets are tested to initialize *k*-means clustering: (i) Handwritten Digits dataset (Dua and Graff 2019) and (ii) NIPS Conference Papers dataset (Perrone et al. 2017). As shown in Figs. 5 and 6, our methods outperform other application-agnostic baselines as well as four baselines for clustering methods.

*Synthetic datasets* We generate a dataset sampled from Gaussian distributions, where  $N = 500$ ,  $d = 5$ , and  $k = 10$ .

*Real-world datasets* Two real-world datasets are tested: (i) Handwritten Digits dataset (Dua and Graff 2019) and (ii) NIPS Conference Papers dataset (Perrone et al. 2017). Handwritten Digits dataset contains 0–9 digit images that can be expressed as  $N = 1797$ ,  $d = 64$ , and  $k = 10$ . NIPS Conference Papers dataset is composed of the papers published from 1987 to 2015. The features of each example are word frequencies, and this dataset can be expressed as  $N = 5811$ ,  $d = 11463$ , and  $k = 20$ . However, without any techniques for reducing the dimensionality, this dataset is hard to apply the



**Fig. 6** Results on initializing  $k$ -means clustering for Handwritten Digits and NIPS Conference Papers datasets

clustering algorithm. We choose 200 dimensions in random when creating the dataset for these experiments, because producing the exact clusters for entire dimensions is not our interest in this paper.

Because the real-world datasets for clustering are difficult to specify truths, we determine truths as class labels for Handwritten Digits dataset (Dua and Graff 2019) and clustering results via Ward hierarchical clustering (Ward 1963) for NIPS Conference Papers dataset (Perrone et al. 2017).

The function of interest in the  $k$ -means clustering setting is the *converged* clustering residual

$$k\text{-means}(\{\mathbf{x}_1, \dots, \mathbf{x}_k\}) = \sum_{i=1}^N \sum_{j=1}^k w_{ij} \|\mathbf{p}_i - \mathbf{c}_j\|_2^2, \tag{28}$$

where  $\{\mathbf{x}_1, \dots, \mathbf{x}_k\}$  is the set of proposed initial cluster centers,  $\{\mathbf{c}_1, \dots, \mathbf{c}_k\}$  is the set of converged cluster centers (Lloyd 1982), and  $w_{ij}$  are softmax values from the pairwise distances. Here, the fact that  $\mathbf{c}_j$  is a function of  $\mathbf{X}$  and  $\mathcal{P}$  is omitted for notational simplicity. The set of converged cluster centers is determined through an iterative strategy which is highly dependent on the initial points  $\mathbf{X}$  to converge to effective centers.

In contrast to  $k$ -means clustering, the GMM estimates parameters of Gaussian distributions and mixing parameters between the distributions. Because it is difficult to minimize negative log-likelihood of the observed data, we fit the GMM using expectation-maximization algorithm (Dempster et al. 1977). Similarly to  $k$ -means clustering, this requires initial guesses  $\mathbf{X}$  to converge to cluster centers  $\{\mathbf{c}_1, \dots, \mathbf{c}_k\}$ .

Table 3 shows convergence qualities and their execution time on  $k$ -means clustering algorithm and GMM for synthetic datasets, and Table 4 represents the qualities and their execution for Handwritten Digits and NIPS Conference Papers datasets. Similar to Table 1, the computational cost increases as  $L$  increases.

**Table 3** Convergence quality and its execution time on  $k$ -means clustering and Gaussian mixture model

$L$	$k$ -means clustering		Gaussian mixture model	
	1-ARI	Time ( $10^3$ s)	1-ARI	Time ( $10^3$ s)
1	$0.051 \pm 0.024$	$3.140 \pm 0.051$	$0.202 \pm 0.011$	$2.364 \pm 0.178$
2	$0.058 \pm 0.030$	$5.034 \pm 0.145$	$0.197 \pm 0.038$	$3.873 \pm 0.148$
5	$0.050 \pm 0.022$	$13.891 \pm 0.327$	$0.194 \pm 0.033$	$13.105 \pm 0.675$
$10(=m)$	$0.073 \pm 0.038$	$45.527 \pm 2.475$	$0.205 \pm 0.030$	$49.134 \pm 1.469$

All settings follow the settings in Fig. 5

**Table 4** Convergence quality and its execution time on  $k$ -means clustering for Handwritten Digits and NIPS Conference Papers datasets

$L$	Handwritten digits		NIPS conference papers	
	1-ARI	Time ( $10^3$ s)	1-ARI	Time ( $10^3$ s)
1	$0.343 \pm 0.007$	$39.347 \pm 6.533$	$0.621 \pm 0.006$	$4.757 \pm 0.138$
2	$0.334 \pm 0.013$	$47.550 \pm 1.832$	$0.620 \pm 0.007$	$9.410 \pm 0.260$
5	$0.340 \pm 0.016$	$177.170 \pm 43.063$	$0.620 \pm 0.007$	$38.521 \pm 1.551$

All settings follow the settings in Fig. 6

#### 4.5 Active nearest neighbor search for point clouds

ModelNet40 dataset (Wu et al. 2015) contains 40 categories of 12,311 3D CAD models. Point cloud representation is obtained by sampling uniformly 1,024 points from the surface of each 3D model using Open3D (Zhou et al. 2018). Nearest neighbor search for point clouds requires large number of Chamfer distance calculations, which is a time-consuming task, in particular when the size of dataset is large. We employ our Bayesian optimization over sets to actively select a candidate whose Chamfer distance from the query is to be computed, while the linear scan requires the calculation of Chamfer distance from the query to every other data in the dataset:

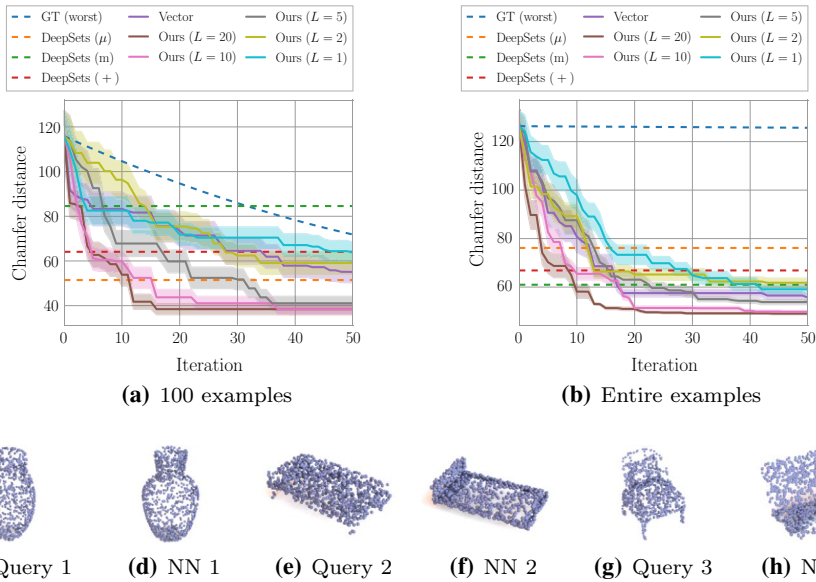
$$d_{\text{Chamfer}}(\mathbf{X}, \mathbf{Y}) = \sum_{\mathbf{x} \in \mathbf{X}} \min_{\mathbf{y} \in \mathbf{Y}} \|\mathbf{x} - \mathbf{y}\|_2 \quad (29)$$

where  $|\mathbf{X}| = |\mathbf{Y}| = m$  and each element in  $\mathbf{X}$  and  $\mathbf{Y}$  is a three-dimensional real vector.

For this experiment, we use one oracle and three additional baselines:

**Ground-truth (worst)** It is the worst case to achieve the ground-truth. The best retrieval case is to find the ground-truth at once, and the usual case would be between the best and the worst.

**DeepSets ( $\mu$ ) / DeepSets (max) / DeepSets (+)** These baselines are implemented to embed point clouds to a single vector using DeepSets (Zaheer et al. 2017), and measure  $l^2$  distance between the embeddings. Because we can access to class information of ModelNet40 dataset (Wu et al. 2015), these neural networks can be trained to match the information. We use three fully-connected layers with batch normalization (Ioffe and Szegedy 2015) as an instance-wise network, and four fully-connected layers with batch normalization as a network after aggregation. ReLU is employed as an activation function. The choice of global aggregation methods determines each baseline: (i) mean aggregation is  $\mu$ ;



**Fig. 7** Nearest neighbor search results on ModelNet40 point clouds. GT and DeepSets (m) indicate the ground-truth and DeepSets (max), respectively. Query and NN pairs are the query and its nearest neighbor examples found by our method

**Table 5** Convergence quality and its execution time on nearest neighbor retrieval for ModelNet40 point clouds

$L$	100 examples		Entire examples	
	Chamfer distance	Time ( $10^3$ s)	Chamfer distance	Time ( $10^3$ s)
1	$63.720 \pm 23.780$	$1.133 \pm 0.049$	$59.095 \pm 10.629$	$1.071 \pm 0.067$
2	$59.079 \pm 25.610$	$1.844 \pm 0.109$	$61.783 \pm 10.000$	$1.835 \pm 0.094$
5	$41.072 \pm 15.807$	$6.568 \pm 0.485$	$53.820 \pm 5.643$	$6.456 \pm 0.372$
10	$38.422 \pm 12.807$	$21.156 \pm 1.360$	$49.798 \pm 3.121$	$23.482 \pm 1.111$
20	$38.422 \pm 12.807$	$104.330 \pm 5.045$	$49.054 \pm 2.614$	$97.179 \pm 3.882$

All settings follow the settings in Fig. 7

(ii) max aggregation is max; and (iii) sum aggregation is +. Because retrieving 1-nearest neighbor with these methods is hard to obtain the nearest neighbor, we choose the nearest neighbor from 3-nearest neighbors to fairly compare with our methods.

Experiments with two different settings were carried out: (i) the size of point clouds is 100; (ii) full-size point clouds (the cardinality is 12,311). In the case of DeepSets, point clouds are embedded into a low-dimensional Euclidean space, so that Euclidean distance is used to search a nearest neighbor (i.e., approximate search). On the other hand, our method actively selects a candidate gradually in the point cloud dataset. The nearest neighbor determined by our method, given the query, is shown in Fig. 7. About 20 iterations of the procedure is required to achieve the better performance, compared to DeepSets (Fig. 7a, b).

## 4.6 Empirical analysis on computational cost

The computational costs from Tables 1, 2, 3, 4 and 5 are presented as a function of  $L$ . These results are measured using a native implementation of set kernels written in Python. As mentioned in Sect. 3, the computational costs follow our expectation, which implies that the complexity for computing a covariance matrix over sets is the major computations in the overall procedure.

## 5 Conclusion

In this paper, we propose the Bayesian optimization method over sets, which takes a set as an input and produces a scalar output. Our method based on GP regression models a surrogate function using set-taking covariance functions, referred to as set kernel. We approximate the set kernel to the efficient positive-definite kernel that is an unbiased estimator of the original set kernel. To find a next set to observe, we employ a constrained acquisition function optimization using the symmetry of the feasible region defined over sets. Moreover, we provide a simple analysis on cumulative regret bounds of our methods. Our experimental results demonstrate our method can be used in some novel applications for Bayesian optimization.

### Declarations

**Conflicts of interest** The authors declare that they have no conflict of interest.

## References

- Arthur, D., & Vassilvitskii, S. (2007). k-means++: The advantages of careful seeding. In *Proceedings of the ACM-SIAM symposium on discrete algorithms (SODA)*, pp. 1027–1035. New Orleans, Louisiana, USA.
- Bergstra, J., Bardenet, R., Bengio, Y., & Kégl, B. (2011). Algorithms for hyper-parameter optimization. In *Advances in neural information processing systems (NeurIPS)*, vol. 24, pp. 2546–2554. Granada, Spain.
- Brochu, E., Cora, V. M., & de Freitas, N. (2010). A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. arXiv preprint [arXiv:1012.2599](https://arxiv.org/abs/1012.2599)
- Buathong, P., Ginsbourger, D., & Kritiyakierne, T. (2020) Kernels over sets of finite sets using RKHS embeddings, with application to Bayesian (combinatorial) optimization. In *Proceedings of the international conference on artificial intelligence and statistics (AISTATS)*, pp. 2731–2741. Virtual.
- Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society B*, 39, 1–38.
- Dua, D., & Graff, C. (2019). UCI machine learning repository. <http://archive.ics.uci.edu/ml>
- Edwards, H., & Storkey, A. (2017). Towards a neural statistician. In *Proceedings of the international conference on learning representations (ICLR)*. Toulon, France.
- Fasshauer, G. E., & McCourt, M. M. (2015). *Kernel-based approximation methods using matlab*. Singapore: World Scientific.

- Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the international conference on machine learning (ICML)*, pp. 1126–1135. Sydney, Australia.
- Frazier, P. I. (2018). A tutorial on Bayesian optimization. arXiv preprint [arXiv:1807.02811](https://arxiv.org/abs/1807.02811)
- Frazier, P. I., & Wang, J. (2016). Bayesian optimization for materials design. In *Information science for materials discovery and design*, pp. 45–75. Springer.
- Garnelo, M., Rosenbaum, D., Maddison, C. J., Ramalho, T., Saxton, D., Shanahan, M., Teh, Y. W., Rezende, D. J., & Eslami, S. M. A. (2018). Conditional neural processes. In *Proceedings of the international conference on machine learning (ICML)*, pp. 1690–1699. Stockholm, Sweden.
- Garnett, R., Osborne, M. A., & Roberts, S. J. (2010). Bayesian optimization for sensor set selection. In *ACM/IEEE international conference on information processing in sensor networks (IPSN)*, pp. 209–219. Stockholm, Sweden.
- Gärtner, T., Flach, P. A., Kowalczyk, A., & Smola, A. J. (2002). Multi-instance kernels. In *Proceedings of the international conference on machine learning (ICML)*, pp. 179–186. Sydney, Australia.
- Genton, M. G. (2001). Classes of kernels for machine learning: A statistics perspective. *Journal of Machine Learning Research*, 2, 299–312.
- González, J., Longworth, J., James, D. C., & Lawrence, N. D. (2014). Bayesian optimization for synthetic gene design. In *Neural information processing systems workshop on Bayesian optimization (Bayes-Opt)*. Montreal, Quebec, Canada.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., & Smola, A. J. (2012). A kernel two-sample test. *Journal of Machine Learning Research*, 13, 723–773.
- Haghanifar, S., McCourt, M., Cheng, B., Wuenschell, J., Ohodnicki, P., & Leu, P. W. (2019). Creating glasswing butterfly-inspired durable antifogging superomniphobic supertransmissive, superclear nanostructured glass through Bayesian learning and optimization. *Materials Horizons*, 6(8), 1632–1642.
- Hansen, N. (2016). The CMA evolution strategy: A tutorial. arXiv preprint [arXiv:1604.00772](https://arxiv.org/abs/1604.00772)
- Haussler, D. (1999). *Convolution kernels on discrete structures*. Tech. rep.: Department of Computer Science, University of California at Santa Cruz.
- Hutter, F., Hoos, H. H., & Leyton-Brown, K. (2011). Sequential model-based optimization for general algorithm configuration. In *Proceedings of the international conference on learning and intelligent optimization (LION)*, pp. 507–523. Rome, Italy.
- Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the international conference on machine learning (ICML)*, pp. 448–456. Lille, France.
- Jones, D. R., Perttunen, C. D., & Stuckman, B. E. (1993). Lipschitzian optimization without the Lipschitz constant. *Journal of Optimization Theory and Applications*, 79(1), 157–181.
- Kandasamy, K., Neiswanger, W., Schneider, J., Póczos, B., & Xing, E.P. (2018). Neural architecture search with Bayesian optimisation and optimal transport. In *Advances in neural information processing systems (NeurIPS)*, vol. 31, pp. 2016–2025. Montreal, Quebec, Canada.
- Kandasamy, K., Schneider, J., & Póczos, B. (2015). High dimensional Bayesian optimisation and bandits via additive models. In *Proceedings of the international conference on machine learning (ICML)*, pp. 295–304. Lille, France.
- Kim, J., McCourt, M., You, T., Kim, S., & Choi, S. (2019). Bayesian optimization over sets. In *International conference on machine learning workshop on automated machine learning (AutoML)*. Long Beach, California, USA.
- Liu, D. C., & Nocedal, J. (1989). On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45(3), 503–528.
- Lloyd, S. P. (1982). Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2), 129–137.
- Muandet, K., Fukumizu, K., Sriperumbudur, B., & Schölkopf, B. (2017). Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends in Machine Learning*, 10(1–2), 1–141.
- Paciorek, C. J., & Schervish, M. J. (2004). Nonstationary covariance functions for Gaussian process regression. In *Advances in neural information processing systems (NeurIPS)*, vol. 17, pp. 273–280. Vancouver, BC, Canada.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12, 2825–2830.
- Perrone, V., Jenkins, P. A., Spano, D., & Teh, Y. W. (2017). Poisson random fields for dynamic feature models. *Journal of Machine Learning Research*, 18, 1–45.
- Rasmussen, C. E., & Williams, C. K. I. (2006). *Gaussian processes for machine learning*. Cambridge: MIT Press.
- Schölkopf, B., & Smola, A. J. (2002). *Learning with kernels*. Cambridge: MIT Press.

- Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., & de Freitas, N. (2016). Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, *104*(1), 148–175.
- Snoek, J., Larochelle, H., & Adams, R. P. (2012). Practical Bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems (NeurIPS)*, vol. 25, pp. 2951–2959. Lake Tahoe, Nevada, USA.
- Srinivas, N., Krause, A., Kakade, S., & Seeger, M. (2010). Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the international conference on machine learning (ICML)*, pp. 1015–1022. Haifa, Israel.
- Wang, Z., Shakibi, B., Jin, L., & de Freitas, N. (2014). Bayesian multi-scale optimistic optimization. In *Proceedings of the international conference on artificial intelligence and statistics (AISTATS)*, pp. 1005–1014. Reykjavik, Iceland.
- Ward, J. H. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, *58*(301), 236–244.
- Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., & Xiao, J. (2015). 3D ShapeNets: A deep representation for volumetric shapes. In *Proceedings of the IEEE international conference on computer vision and pattern recognition (CVPR)*, pp. 1912–1920. Boston, MA, USA.
- Zaheer, M., Kottur, S., Ravanbakhsh, S., Poczos, B., Salakhutdinov, R. R., & Smola, A. J. (2017). Deep sets. In *Advances in neural information processing systems (NeurIPS)*, vol. 30, pp. 3391–3401. Long Beach, CA, USA.
- Zhou, Q., Park, J., & Koltun, V. (2018). Open3D: A modern library for 3D data processing. arXiv preprint [arXiv:1801.09847](https://arxiv.org/abs/1801.09847)
- Zoph, B., & Le, Q. V. (2017). Neural architecture search with reinforcement learning. In *Proceedings of the international conference on learning representations (ICLR)*. Toulon, France.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.