

Bayesian Optimization over Permutation Spaces

Aryan Deshwal, Syrine Belakaria, Janardhan Rao Doppa, Dae Hyun Kim

School of EECS, Washington State University
{aryan.deshwal, syrine.belakaria, jana.doppa, daehyun.kim}@wsu.edu

Abstract

Optimizing expensive to evaluate black-box functions over an input space consisting of all permutations of d objects is an important problem with many real-world applications. For example, placement of functional blocks in hardware design to optimize performance via simulations. The overall goal is to minimize the number of function evaluations to find high-performing permutations. The key challenge in solving this problem using the Bayesian optimization (BO) framework is to trade-off the complexity of statistical model and tractability of acquisition function optimization. In this paper, we propose and evaluate two algorithms for **BO over Permutation Spaces (BOPS)**. First, BOPS-T employs Gaussian process (GP) surrogate model with Kendall kernels and a Tractable acquisition function optimization approach based on Thompson sampling to select the sequence of permutations for evaluation. Second, BOPS-H employs GP surrogate model with Mallow kernels and a Heuristic search approach to optimize expected improvement acquisition function. We theoretically analyze the performance of BOPS-T to show that their regret grows sub-linearly. Our experiments on multiple synthetic and real-world benchmarks show that both BOPS-T and BOPS-H perform better than the state-of-the-art BO algorithm for combinatorial spaces. To drive future research on this important problem, we make new resources and real-world benchmarks available to the community.

1 Introduction

Optimizing expensive black-box functions is a common problem in many science and engineering applications (Frazier and Wang 2016; Yang, Wu, and Arnold 2019; Zhou et al. 2020; Deshwal, Simon, and Doppa 2021). An important class of such problems involve optimizing over the space of all permutations of a set of objects. For example, in the design of integrated circuits (ICs), we need to find placement of functional blocks to optimize performance via expensive simulations. Another example is to find a schedule for the jobs in additive manufacturing (parts can be built layer-by-layer) to optimize throughput via simulations. Bayesian optimization (BO) (Shahriari et al. 2016) has proven to be a successful approach for optimizing black-box functions in a sample-efficient manner. The key idea is to learn a surrogate statistical model, e.g., Gaussian process, and use it to select

the sequence of inputs for function evaluation to uncover high-performing inputs.

A large body of the BO literature focuses on continuous input spaces (Greenhill et al. 2020) with few recent works (Oh et al. 2019; Baptista and Poloczek 2018; Dadkhahi et al. 2020; Doppa 2021) tackling the challenging setting of input spaces over discrete structures, e.g., sets, sequences, and graphs. Unlike continuous spaces, discrete spaces come with many unique challenges such as difficulty of defining a general representation, non-smoothness, etc. which require specialized treatment of different types of combinatorial structures (BakIr et al. 2007). In this paper, we consider the understudied space of all *permutations* of d objects, which is equivalently characterized by the non-Abelian Symmetric group \mathcal{S}_d of cardinality $d!$, consisting of all bijections from a set to itself. In contrast, all of the existing work on combinatorial BO is focused on the input space with *categorical/binary* variables which corresponds to the direct sums of abelian cyclic groups $\mathbb{Z}/c\mathbb{Z}$ of cardinality c^d , where each of the d variables can take one of the c categories. Combinatorial BO methods such as COMBO (Oh et al. 2019) cannot exploit the special characteristics of permutation spaces. For example, they account for d^d space $\gg d!$ permutation space.

The key challenge to devise effective solutions for **BO over Permutation Spaces (BOPS)** is to trade-off the complexity of statistical model and tractability of search to select permutations with high utility/acquisition values (e.g., expected improvement). We propose and evaluate two algorithms for BOPS by addressing this challenge. First, BOPS-T uses a simpler statistical model in the form of GP with Kendall kernels (Jiao and Vert 2015). By considering the weight space view of the GP model and employing Thompson sampling as the acquisition strategy, we show that the acquisition function optimization is a Quadratic Assignment problem which can be solved using Tractable Semi-definite programming relaxation based solvers. Second, BOPS-H uses a complex statistical model in the form of GP with Mallow kernels (Jiao and Vert 2015; Mania et al. 2018). We employ expected improvement as the acquisition strategy and perform Heuristic search (local search with multiple restarts) to select the permutations for function evaluation. We analyze the theoretical properties of BOPS-T and in terms of regret bounds and show that it achieves a sublinear (time) regret. Our comprehensive experiments on both syn-

thetic and real-world benchmarks show that both BOPS-T and BOPS-H perform better than COMBO (Oh et al. 2019).

Contributions. The main contributions of this paper are:

- Two BO algorithms for permutation spaces, namely, BOPS-T and BOPS-H, that make varying trade-offs between the complexity of statistical model and tractability of search for selecting permutations for evaluation.
- We theoretically analyze BOPS-T and show that it achieves sublinear (in time) regret bounds of $\mathcal{O}^*(d^{3/2}\sqrt{T})$, where d and T refers to dimensionality and time (no. of iterations) respectively and \mathcal{O}^* denotes upto logarithmic factors.
- We evaluate the efficacy of our algorithms and compare with the state-of-the-art COMBO algorithm on multiple synthetic and real-world benchmarks.
- We make new resources and benchmarks based on important real-world problems available to the BO community to drive future research on this under-studied problem. The resources and source code of BOPS algorithms are available at <https://github.com/aryandeshwal/BOPS/>.

2 Problem Formulation

In this paper, we consider optimization problems with the input space consisting of all permutations over d objects. Given $[1, d] := \{1, 2, \dots, d\}$, indexing the d objects, a permutation is defined as a bijective mapping $\pi : [1, d] \mapsto [1, d]$. The set of all permutations along with the composition binary operation $((\pi_1 \circ \pi_2)(x) = \pi_1(\pi_2(x)))$ is known as the Symmetric group \mathcal{S}_d which has a cardinality $|\mathcal{S}_d| = d!$.

Let $f : \mathcal{S}_d \mapsto \mathbb{R}$ be a black-box objective function that is expensive to evaluate. Our goal is to optimize f while *minimizing the number of function evaluations*:

$$\pi^* = \arg \min_{\pi \in \mathcal{S}_d} f(\pi) \quad (2.1)$$

For a concrete example problem, consider the domain of design and optimization of integrated circuits (ICs). There are many applications in IC design, where we need to optimize over permutations of functional blocks of different granularity (small cells to processing cores). Some example objectives include performance and manufacturing cost. We need to perform expensive computational simulations to evaluate each candidate permutation.

We solve this problem using the Bayesian optimization (BO) framework. A BO algorithm is composed of two main components: 1) a *surrogate statistical model* built using past evaluations of the expensive black-box objective function; and 2) an *acquisition function* to capture the expected utility of evaluating a new input, which is optimized at each BO iteration to find the next best input for evaluation. Algorithm 1 shows the generic BOPS approach. We provide concrete instantiations for these two components by handling the unique challenges of permutation spaces. In each BO iteration, we select the most promising permutation for evaluation (line 3) and update the statistical models using all the training examples (line 6). The best found permutation is returned at the end of maximum BO iterations (line 8).

Algorithm 1: Bayesian Optimization over Permutation Spaces (BOPS)

requires: black-box objective $f(\pi)$, Gaussian process $GP(0, k)$ with Kernel over permutations k , Acquisition function AF

- 1: $D_0 \leftarrow$ small initial data; and $GP_0 \leftarrow GP(\mu_{D_0}, k_{D_0})$
 - 2: **for** $j=1, 2, \dots$ **do**
 - 3: Acquisition function optimization to select the next permutation: $\pi_j = \arg \min_{\pi \in \mathcal{S}_d} AF(\pi)$
 - 4: Evaluate the selected permutation π_j to get $f(\pi_j)$
 - 5: Aggregate training data: $D_j \leftarrow D_{j-1} \cup \{\pi_j, f(\pi_j)\}$
 - 6: Update GP posterior: $GP_j \leftarrow GP(\mu_{D_j}, k_{D_j})$
 - 7: **end for**
 - 8: **return** $\pi_{best} = \arg \min\{f(\pi_1), f(\pi_2) \dots\}$
-

3 Related Work

Bayesian optimization (BO) methods for input spaces over discrete structures can be broadly classified into two categories based on the amount of available training data: *large* and *small* data settings. Our paper focuses on the small-data setting due to the unavailability of large training sets and the objective function evaluations being expensive.

BO methods for large-data setting. Most of these methods employ the large set of training examples to learn embeddings of discrete structures in a latent space and perform continuous BO in the latent space (Gómez-Bombarelli et al. 2018; Lu et al. 2018; Eissman et al. 2018; Kajino 2019). This setting is especially relevant for domains such as molecular design (Griffiths and Hernández-Lobato 2017; Hernández-Lobato et al. 2017) and biological sequence design (Angermueller et al. 2020, 2019).

BO methods for small-data setting. Prior work in the small-data setting consider input spaces over binary/categorical structures: $x_1 \times x_2 \dots \times x_d$ space, where each x_i can take all possible values from a set of categories $\{0, 1, \dots, c-1\}$, where $c=2$ and $c > 2$ for binary and categorical variables respectively. This results in a search space of size c^d , which has a different structure than our considered setting of the space of all permutations of a set of d objects. Prior methods in this setting can be classified based on their choice of surrogate models (linear models (Baptista and Poloczek 2018; Dadkhahi et al. 2020), random forests (Hutter, Hoos, and Leyton-Brown 2011), Gaussian processes (Oh et al. 2019; Deshwal and Doppa 2021)) or acquisition function optimization approaches (heuristic search (Hutter, Hoos, and Leyton-Brown 2011; Oh et al. 2019; Dadkhahi et al. 2020; Deshwal and Doppa 2021), principled mathematical optimization (Baptista and Poloczek 2018; Deshwal, Belakaria, and Doppa 2020, 2021)), and combination thereof (Garrido-Merchán and Hernández-Lobato 2020; Luong et al. 2019; Deshwal et al. 2020; Swersky et al. 2020).

Gaussian process models for discrete structures. COMBO (Oh et al. 2019) is the state-of-the-art method for small-data setting. It employs Gaussian process (GP) with diffusion kernels (Kondor and Lafferty 2002) defined over a combinatorial graph representation of the input space. However, COMBO’s graph representation cannot

exploit the special characteristics of permutation spaces. For example, for a space of d objects, it requires d subgraphs each with at least d nodes to account for d^d combinatorial space which is unnecessarily large for our usage. Another line of work considers GP models with kernels for specific structures including sets (double sum/embedding kernels (Buathong, Ginsbourger, and Krityakierne 2020; Kim et al. 2021)), molecular graphs (optimal transport based kernel (Korovina et al. 2020)), and strings (string kernels (Moss et al. 2020; Lodhi et al. 2002)). We employ Kendall kernels and Mallow kernels (Jiao and Vert 2015; Mania et al. 2018) for permutation spaces which were shown to have superior performance on permutation based classification tasks.

4 BO Algorithms for Permutation Spaces

In this section, we provide two algorithms for BO over permutation spaces that make varying trade-offs between the complexity of statistical model and tractability of acquisition function optimization. First, BOPS-T employs a simple statistical model with an efficient Semi-definite programming (SDP) relaxation based optimization method. Second, BOPS-H employs a complex statistical model and performs heuristic search for optimizing the acquisition function. We employ Gaussian processes (GPs) (Rasmussen and Williams 2006) as the surrogate model in both algorithms. GPs are effective statistical models commonly used for BO as they provide a principled framework for uncertainty estimation. They are fully characterized by a kernel k (Kanagawa et al. 2018) which intuitively captures the similarity between two candidate inputs from the same input space.

4.1 BOPS-T Algorithm

Surrogate model. The similarity between any two permutations (π, π') can be naturally defined by considering the number of pairs of objects ordered in the same way or in opposite ways. This is captured by the notion of the number of discordant pairs $n_d(\pi, \pi')$, also known as Kendall-tau distance (Kendall 1938). $n_d(\pi, \pi')$ counts the number of pairs of objects ordered oppositely by π and π' as defined below:

$$n_d(\pi, \pi') = \sum_{i < j} [1_{\pi(i) > \pi(j)} 1_{\pi'(i) < \pi'(j)} + 1_{\pi(i) < \pi(j)} 1_{\pi'(i) > \pi'(j)}] \quad (4.1)$$

A related notion of concordant pairs $n_c(\pi, \pi')$ counts the number of object pairs ordered similarly by π and π' :

$$n_c(\pi, \pi') = \binom{d}{2} - n_d(\pi, \pi') \quad (4.2)$$

Kendall kernels (Jiao and Vert 2015) are positive-definite kernels defined over permutations using the notion of discordant and concordant pairs as follows:

$$k(\pi, \pi') = \frac{n_c(\pi, \pi') - n_d(\pi, \pi')}{\binom{d}{2}} \quad (4.3)$$

Because of their proven effectiveness over permutations (Jiao and Vert 2015, 2018), we propose using Kendall kernels with GPs as surrogate model in our BOPS-T algorithm.

For our surrogate model, we consider the weight-space formulation of the GP. This weight-space formulation is essential for the SDP based acquisition function optimization approach described in the next section. In the *weight-space* view, we can reason about GPs as a weighted sum of basis functions $\phi = \{\phi_i(\cdot)\}$, i.e.,

$$w^T \phi(\cdot); \quad w \sim N(0, I) \quad (4.4)$$

where $N(\cdot)$ represents multi-variate Gaussian distribution and I is the identity matrix. Every kernel has a canonical feature map (as per the Moore-Aronszajn theorem (Aronszajn 1950)) $\phi : \mathcal{S}_d \mapsto H_k$, H_k being its associated Reproducing Kernel Hilbert Space (RKHS), that is employed as the basis function in 4.4. The feature map expression for Kendall kernel (constructed by (Jiao and Vert 2015)) is given below:

$$\phi(\pi) = \left\{ \sqrt{\binom{d}{2}^{-1}} (1_{\pi(i) > \pi(j)} - 1_{\pi(i) < \pi(j)}) \right\}_{(1 \leq i < j \leq d)} \quad (4.5)$$

Acquisition function and optimizer. In order to sequentially select the next permutation for evaluation guided by the learned surrogate model, we employ Thompson sampling as our acquisition function. Thompson sampling is a powerful, practitioner-friendly, and parameter-free approach for appropriately balancing the exploration vs. exploitation dilemma (Russo and Van Roy 2014) in sequential bandit optimization. The key idea is to sample a function from the surrogate model's posterior and select its optimizer as the next permutation for evaluation. In the weight-space view of GPs, this corresponds to sampling a weight vector \hat{w} from its posterior and solving the following optimization problem:

$$\pi_{next} = \arg \min_{\pi \in \mathcal{S}_d} \hat{w}^T \phi(\pi) \quad (4.6)$$

It should be noted that the sampled weight vector \hat{w} is an exact function defined by GP (with Kendall kernel) over permutation spaces and has no approximation error when compared to the function space approach. This is in contrast to the common practice of using Thompson sampling over continuous spaces, where random Fourier features based weight-space representation of GPs is used which inevitably results in approximation error because of sampling a finite number of features from an infinite feature space.

We now show that the above acquisition function optimization problem (4.6) is a Quadratic Assignment Problem (QAP) (Burkard et al. 1998). To observe that, the objective in 4.6 is written in an equivalent form in terms of P_d , the set of all possible permutation matrices P of size $d \times d$, as follows:

$$\min_{P \in P_d} Tr(WPAP^T) \quad (4.7)$$

where Tr is the matrix trace operation and A is a $d \times d$ matrix defined as follows:

$$A = \begin{cases} 1 & \text{if } i < j \\ -1 & \text{if } i > j \\ 0 & \text{if } i == j \end{cases} \quad \forall i, j \in [1, d]$$

and W is another $d \times d$ matrix given as follows:

$$W = \begin{cases} w_{\frac{(i-1)}{2}(2d-i)+(j-i)} & \text{if } i < j \\ 0 & \text{if } i \geq j \end{cases} \quad \forall i, j \in [1, d]$$

Concretely, the equivalence of objectives in 4.6 and 4.7 can be seen as follows:

$$Tr(WPAP^T) = \sum_{i=1}^d (WPAP^T)_{ii} \quad (4.8)$$

Equation 4.8 is the definition of the trace of a matrix. Now, considering each entry $(WPAP^T)_{ii}$ in 4.8:

$$(WPAP^T)_{ii} = \sum_{j=1}^d W_{ij} \cdot (PAP^T)_{ji} \quad (4.9)$$

$$= \sum_{j>i}^d w_{\frac{(i-1)}{2}(2d-i)+(j-i)} \cdot (PAP^T)_{ji} \quad (4.10)$$

$$= \sum_{j>i}^d w_{\frac{(i-1)}{2}(2d-i)+(j-i)} \cdot A_{\pi(j)\pi(i)} \quad (4.11)$$

where 4.10 follows from the definition of W and 4.11 follows from the fact that pre-multiplying (post-multiplying) by a permutation matrix permutes the rows (columns) of A . Using 4.11 in 4.8:

$$Tr(WPAP^T) = \sum_{i=1}^d \sum_{j>i}^d w_{\frac{(i-1)}{2}(2d-i)+(j-i)} \cdot A_{\pi(j)\pi(i)} \quad (4.12)$$

By noting that $A_{\pi(j)\pi(i)}$ is exactly the feature map in 4.5 (upto multiplication by a constant $\sqrt{\binom{d}{2}^{-1}}$ which doesn't change the optimal solution), the equivalence between 4.6 and 4.7 is established.

Although, in general, Quadratic assignment problem is NP-hard (Sahni and Gonzalez 1976), we leverage existing Semi-definite programming (SDP) based strong relaxations (Zhao et al. 1998) to obtain good approximate solutions to the acquisition function optimization problem. Using the invariance of the trace under cyclic permutations and vectorization identity ($vec(APW) = (W^T \otimes A)vec(P)$), objective in 4.7 is standardized as:

$$\begin{aligned} \min_{P, Q} ((W^T \otimes A)Q) & \quad (4.13) \\ P \in P_n & \\ Q = vec(P)vec(P)^T & \end{aligned}$$

where $vec(P)$ is the column-wise vectorization of P . We leverage the clique-based SDP relaxation approach of (Ferreira, Khoo, and Singer 2018) which can exploit matrix sparsity (e.g., zeros in the upper-triangular matrix W) for solving 4.13. The key idea is to enforce semi-definiteness only over groups of Q 's entries (i.e., cliques) to get a relaxation that can be solved using fast and accurate algorithms.

4.2 BOPS-H Algorithm

Surrogate model. We propose to employ Mallows kernel which plays a role on the symmetric group \mathcal{S}_d similar to the Gaussian (RBF) kernel on the Euclidean space. Given a pair of permutations π and π' , the Mallows kernel is defined as the exponentiated negative of the number of discordant pairs $n_d(\pi, \pi')$ between π and π' i.e.

$$k_m \pi, \pi' = \exp(-ln_d(\pi, \pi')) \quad (4.14)$$

where $l \geq 0$ is a hyper-parameter of the kernel similar to the length-scale hyper-parameter of the Gaussian kernels on Euclidean space. A key measure of the expressivity of a kernel is based on a property called *universality* which captures the notion of whether the RKHS of the kernel is rich enough to approximate any function on a given input space arbitrary well. It was recently shown (Mania et al. 2018) that Mallows kernel is *universal* over the space of permutations in contrast to the Kendall kernel discussed in the previous section. Therefore, Mallows kernels are more powerful than Kendall kernels and allows us to capture richer structure in permutations when used to learn GP based surrogate models. Indeed, our experiments also demonstrate empirically the superior modeling capability of Mallows Kernel.

Acquisition function and optimizer. Unlike Kendall kernel, the feature space of Mallows Kernel is exponentially large (Mania et al. 2018) making it practically inefficient to sample functions from the GP posterior (in the weight-space style as described earlier). Therefore, we propose to employ expected improvement (EI) as our acquisition function. The additional complexity of GP based statistical model with Mallows kernel makes the acquisition function optimization problem $\pi_{next} = arg \min_{\pi \in \mathcal{S}_d} AF(\pi)$ is intractable for EI. Therefore, we propose to perform **H**euristic search in the form of local search with multiple restarts that has been shown to be very effective in practice for solving combinatorial optimization problems. To search over only valid permutations $\pi \in \mathcal{S}_d$, at each local search step, we consider only those neighbors which are permutations of the current state. Otherwise, we will be searching over a huge combinatorial space with both valid (permutations) and invalid structures (non-permutations) as done by COMBO: may not result in producing a permutation from its acquisition function optimization procedure. Indeed, we observed this behavior in our experiments with COMBO. We use the modified local search procedure over permutations for both COMBO and our BOPS-H algorithm in experiments.

5 Theoretical Analysis for BOPS-T

In this section, we analyze the theoretical properties of our BOPS-T algorithm in terms of regret metric (Srinivas et al. 2009), which is a commonly used measure for analyzing BO algorithms. Note that there is no prior regret bound analysis for BO algorithms for EI even in continuous spaces. Hence, we leave the analysis of BOPS-H algorithm for future work. Let simple regret R be defined as follows:

$$R = \sum_{t=1}^T (f(\pi_t) - f(\pi^*)) \quad (5.1)$$

where π_t is the permutation picked by the BO algorithm at time (iteration) t . In our case of using Thompson sampling as an acquisition function, it is natural to consider the expected form of this regret (Russo and Van Roy 2014) where the expectation is taken over the distribution of functions as given by the GP prior with Kendall kernel. We analyze this expected form of regret, also known as Bayesian regret:

$$\mathcal{BR} = \sum_{t=1}^T \mathbb{E}(f(\pi_t) - f(\pi^*)) \quad (5.2)$$

where the expectation is over the distribution of functions $f \sim GP(0, k)$. The below theorem bounds the Bayesian regret of our BOPS-T algorithm:

Theorem 5.1 *Let $f \sim GP(0, k)$ with Kendall kernel k (4.3), the Bayesian regret of the BOPS-T algorithm after T observations $y_i = f(\pi_i) + \epsilon_i, i \in \{1, 2, \dots, T\}$ with ϵ_i being Gaussian distributed i.i.d. noise $\epsilon_i \sim N(0, \sigma^2)$ is : $\mathcal{BR} = \mathcal{O}^*(d^{3/2}\sqrt{T})$, where \mathcal{O}^* denotes upto log factors.*

Proof. The key quantity in bounding the regret of Bayesian optimization with Gaussian processes (also known as GP bandits) is an information-theoretic quantity called as *maximum information gain* γ_T (Srinivas et al. 2009) that depends on the kernel k and intuitively captures the maximum information that can be gained about f after T observations, i.e.,

$$\gamma_T = \max_{A \subset \mathcal{S}_d, |A|=T} I(\mathbf{y}_A; f) \quad (5.3)$$

where I is the mutual information and A is a subset of permutations with corresponding function evaluations \mathbf{y}_A .

(Russo and Van Roy 2014) proved the Bayesian regret for Thompson sampling by characterizing it in terms of upper confidence bound based results from (Srinivas et al. 2009):

Proposition 1 (*Proposition 5 (Russo and Van Roy 2014)*). *If $|X| < \infty$, $\{f(x) : x \in X\}$ follows a multivariate Gaussian distribution with marginal variances bounded by 1, the Bayesian regret for Thompson sampling based bandit policy is given as:*

$$\mathcal{BR} = 1 + 2\sqrt{T\gamma_T \ln(1 + \sigma^2)^{-1} \ln\left(\frac{(T^2 + 1)|X|}{\sqrt{2\pi}}\right)} \quad (5.4)$$

where X is the action space.

This proposition is directly applicable in our setting because the action space, being the cardinality of the symmetric group \mathcal{S}_d , is finite (i.e., $|\mathcal{S}_d| = d!$) and the function $\{f(\pi) : \pi \in \mathcal{S}_d\}$ follows a multivariate Gaussian distribution (by definition of Gaussian process with Kendall kernel). We compute the specific terms in the right-hand side of 5.4 that are applicable in our setting to prove the regret bound.

The maximum information gain for kernels with finite feature maps can be computed in the weight-space form (Sec 4.1) as a special case of linear kernel (Srinivas et al. 2009).

$$\gamma_T \leq C \log |I + \sigma^{-2}K| \quad (5.5)$$

where $C = 1/2 \cdot (1 - 1/e)^{-1}$ is a constant, K is a $T \times T$ matrix with each entry $K_{ij} = k(\pi_i, \pi_j)$. As per kernel trick,

$$K = \Phi^T \Phi \quad (5.6)$$

where Φ is a matrix with $\Sigma^{1/2}\phi(\pi_i), i \in \{1, 2, \dots, T\}$ as the columns (4.5). Therefore,

$$\gamma_T \leq C \ln |I + \sigma^{-2}\Phi^T \Phi| \quad (5.7)$$

By Schur's complement:

$$\gamma_T \leq C \ln |I + \sigma^{-2}\Phi^T \Phi| \leq C \ln |I + \sigma^{-2}\Phi \Phi^T| \quad (5.8)$$

By Hadamard's inequality:

$$\gamma_T \leq C \ln |I + \sigma^{-2}\Phi \Phi^T| \quad (5.9)$$

$$\leq C \sum_{i=1}^{\binom{d}{2}} \ln(1 + \sigma^{-2}\lambda_i) \quad (5.10)$$

where $\{\lambda_1, \lambda_2, \dots\}$ is the eigenvalue set of the matrix $\Phi \Phi^T$.

By Gershgorin circle theorem (Varga 2010), all the eigenvalues of a matrix is upper bounded by the maximum absolute sum of rows, i.e. $\lambda_i \leq d^2 T$ with the assumption that $\|\Sigma^{1/2}\phi(\pi)\| \leq 1$.

$$\gamma_T = O(d^2 \ln(d^2 T)) \quad (5.11)$$

Now, using Stirling's approximation, we can bound the $\ln(|X|)$ term in 5.4, where $|X| = |\mathcal{S}_d|$ in our case:

$$\ln(|\mathcal{S}_d|) = O(d \ln d) \quad (5.12)$$

Plugging 5.11 and 5.12 in 5.4 and ignoring constants, we get the following expression:

$$\mathcal{BR} = O(\sqrt{T d^2 \ln d^2 T (\ln T^2 + d \ln d)}) \quad (5.13)$$

$$\mathcal{BR} = O(\sqrt{(T d^2 \ln d^2 T \ln T^2 + T d^3 \ln d^2 T \ln d)}) \quad (5.14)$$

$$\mathcal{BR} = \mathcal{O}^*(d^{3/2}\sqrt{T}) \quad (5.15)$$

Hence, ignoring log factors, our proposed BOPS-T algorithm achieves sublinear (time) regret.

6 Experiments and Results

In this section, we describe the benchmarks and experimental setup followed by results and discussion.

6.1 Benchmarks

We employ diverse and challenging benchmarks for black-box optimization over permutations for our experiments. We have the following two synthetic benchmarks.

1) Quadratic assignment problem (QAP). QAPLIB (Burkard, Karisch, and Rendl 1997) is a popular library that contains multiple QAP instances. Each QAP instance contains a cost matrix (A) and distance matrix (B) sized $n \times n$, where n is the number of input dimensions. The goal is to find the best permutation that minimizes the quadratic assignment objective $Tr(APBP^T)$, where P is an $n \times n$ permutation matrix. We use input space with $n = 15$ dimensions in our experiments.

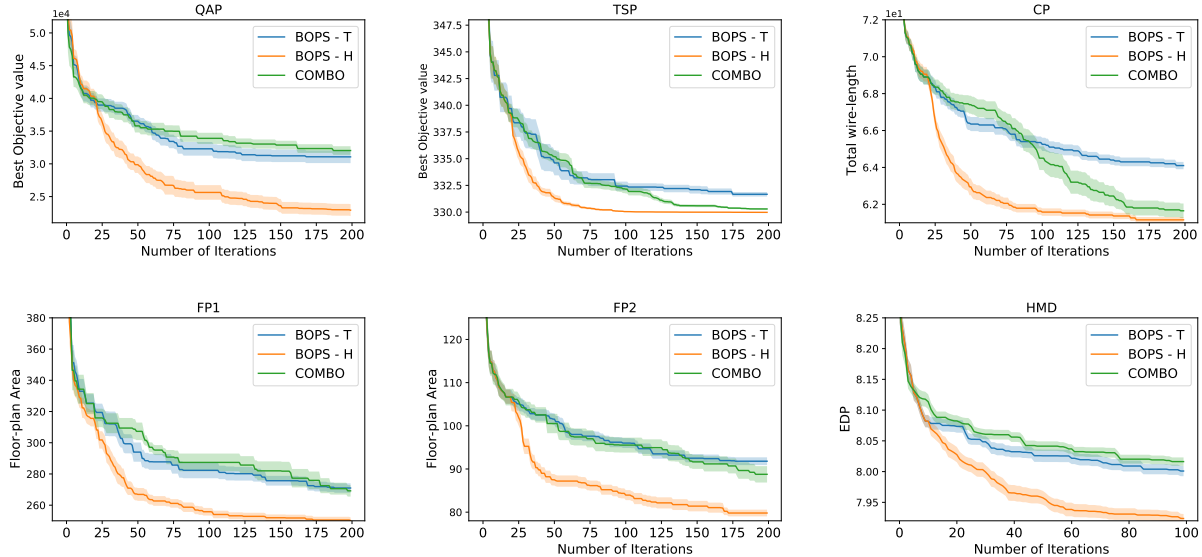


Figure 1: Results comparing BOPS-T, BOPS-H, and COMBO (best objective function value vs. number of BO iterations) on both synthetic and real-world benchmarks: (Top row) QAP, TSP, CP; and (Bottom row) FP1, FP2, and HMD.

2) Traveling salesman problem (TSP). TSP problems are derived from low-dimensional variants of the printed circuit board (PCB) problems from the TSPLIB library (Reinelt 1995). The overall goal is to find the route of drilling holes in the PCB that minimizes the time taken to complete the job. We use input space with $d = 10$ dimensions from the data provided in the library.

We perform experiments on three important real-world applications from the domain of computer-aided design of integrated circuits (ICs). These applications are characterized by permutations over functional blocks at different levels of granularity that arise in different stages of design and optimization of ICs. Importantly, even tiny improvements in solution has huge impact (e.g., improved performance over the lifespan of the IC or reduced cost for manufacturing large samples of the same IC). A big challenge in the combinatorial BO literature is the availability of challenging real-world problems to evaluate new approaches. Hence, we provide our three real-world benchmarks as a new resource to allow rapid development of the field.

3) Floor planning (FP). We are given k rectangular blocks with varying width and height, where each block represents a functional module performing certain task. Each placement of the given blocks is called a *floor-plan*. Our goal is to find the floor plan that minimizes the manufacturing cost per chip. We use two variants of this benchmark with 10 blocks (FP1 and FP2) that differ in the functionality of the blocks.

4) Cell placement (CP). We are given 10 rectangular cells with same height and a netlist that contains the connection information among the cells. The goal is to place the 10 rectangular cells for optimizing the performance of the circuit. Intuitively, shorter nets have shorter delays, so placements with shorter wire-length will result in higher performance.

5) Heterogeneous manycore design (HMD). This is a

manycore architecture optimization problem from the ro-dinia benchmark (Che et al. 2009). We are given 16 cores of three types: 2 CPUs, 10 GPUs, and 4 memory units. They are connected by a mesh network (each core is connected to its four neighboring cores) to facilitate data transfer. The goal is to place the given 16 cores to optimize the energy delay product (EDP) objective that captures both latency and energy, two key attributes of a manycore chip.

6.2 Experimental Setup

Configuration of algorithms. We compare our proposed BOPS-T and BOPS-H algorithms with the state-of-the-art combinatorial BO algorithm COMBO (Oh et al. 2019). COMBO employs a diffusion kernel based GP surrogate model and optimizes expected improvement acquisition function using local search with restarts to select inputs for evaluation. Each local search step considers all neighbors of the current structure in the combinatorial graph (i.e., structures with Hamming distance one). We modify COMBO’s local search procedure (<https://github.com/QUVA-Lab/COMBO>) to consider only those neighbors which are permutations of the current state thereby helping COMBO to avoid searching a large combinatorial space with huge number of invalid structures (non-permutations).

We used the SDP relaxation based QAP solver code from (<https://github.com/fsbravo/csdp>) for implementing BOPS-T. BOPS-H is built using popular GPyTorch (Gardner et al. 2018) and BoTorch (Balandat et al. 2020) libraries. We used 10 restarts for local search based EI optimization for BOPS-H. BOPS-T, BOPS-H, and COMBO are initialized with the same 20 random permutations in each experiment.

Evaluation metric. We plot the objective function value of the best permutation over different BO iterations. Each experiment is repeated 20 times and we plot the mean of the

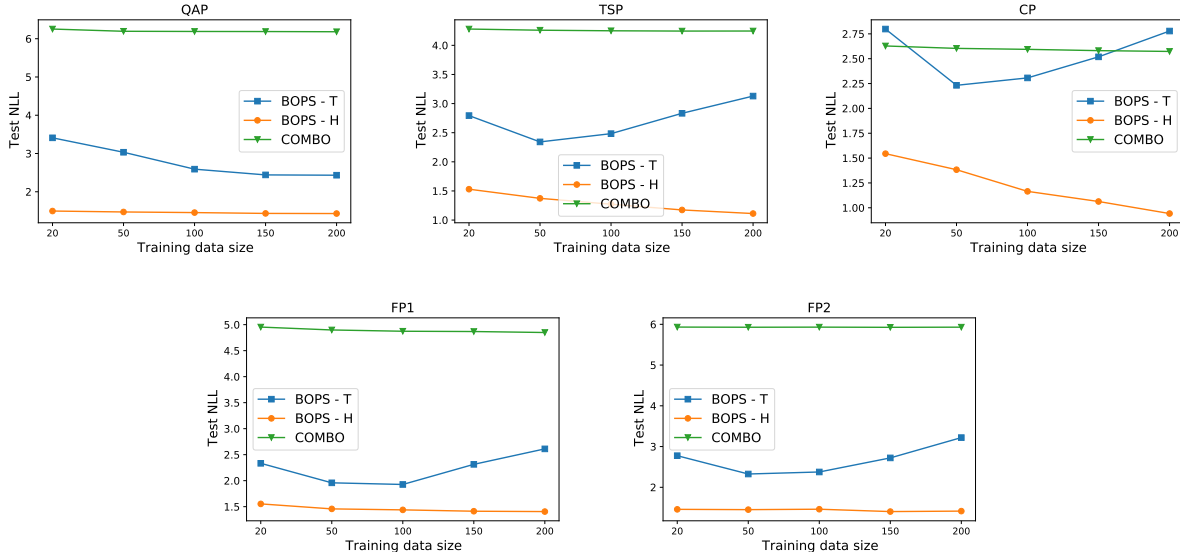


Figure 2: Results comparing the three surrogate models of BOPS-T, BOPS-H and COMBO on negative log-likelihood (NLL) metric computed on a test set on five benchmarks: (Top row) QAP, TSP, CP; and (Bottom row) FP1, FP2.

best objective value plus and minus the standard error.

6.3 Results and Discussion

In this section, we present and discuss our experimental results along different dimensions.

Figure 1 shows the results for BO performance (best objective value vs. number of function evaluations / BO iterations) of BOPS-T, BOPS-H, and COMBO on all six benchmarks. Below we discuss these results in detail.

BOPS-T vs. BOPS-H. Recall that BOPS-T and BOPS-H makes varying trade-offs between the complexity of statistical model and tractability of acquisition function optimization: BOPS-T uses simple model and tractable search; and BOPS-H employs complex model and heuristic search. From Figure 1, we can observe that BOPS-H performs significantly better than BOPS-T on all six benchmarks.

BOPS vs. COMBO. From the results shown in Figure 1, we make the following observations: 1) BOPS-H performs significantly better than both BOPS-T and COMBO on all six benchmarks; and 2) BOPS-T is comparable or slightly better than COMBO on all benchmarks except TSP and CP.

We hypothesize that the performance of different BO algorithms, namely, BOPS-H, BOPS-T, and COMBO is proportional to the quality of their surrogate models in terms of making predictions on unknown permutations and their uncertainty quantification ability. To verify this hypothesis, we compare the three surrogate models quantitatively in terms of their performance on the log-likelihood metric.

Comparison of surrogate models. We compare the three surrogate models on the log-likelihood metric (Murphy 2012) because it captures both the prediction and uncertainty quantification of a model which are essential for the effectiveness of BO. We plot the negative log-likelihood (NLL) of the three surrogate models on a testing set of 50 instances

as a function of the increasing size of training data. Each experiment is replicated with 10 different training sets and each method is evaluated using the median of the NLL metric on 10 different test sets of 50 permutations each. Figure 2 shows the results on all benchmarks except HMD. We do not show results on HMD since each function evaluation is much more expensive when compared to all other benchmarks, and we are generating multiple replications of the training and testing sets ($10 \times 10 = 100$ runs). We make the following observations from Figure 2: 1) BOPS-H shows the best performance among the three surrogate models; 2) BOPS-T does better than COMBO on all benchmarks other than cell-placement. Since both COMBO and BOPS-H employ the same acquisition function (EI) and optimizer (local search), it is evident that the gains in the BO performance comes from the superior surrogate model of BOPS-H.

7 Conclusions

We proposed and evaluated two effective Bayesian optimization algorithms with varying trade-offs for optimizing expensive black-box functions over the challenging input space of permutations. The results point to a key conclusion that it is important to use an appropriate model that exploits the specific structure of permutation spaces, which is different than the generic combinatorial space over categorical variables. We characterized the importance of this problem setting by describing three important real-world applications from the domain of computer-aided design of integrated circuits. Furthermore, we make all these benchmarks available to drive future research in this problem space. Future work includes studying extensions to handle high-dimensional spaces (Oh, Gavves, and Welling 2018) and multiple objectives (Belakaria, Deshwal, and Doppa 2019).

Acknowledgments. This research is supported in part by NSF grants IIS-1845922, OAC-1910213, and SII-2030159.

References

- Angermueller, C.; Belanger, D.; Gane, A.; Mariet, Z.; Dohan, D.; Murphy, K.; Colwell, L.; and Sculley, D. 2020. Population-Based Black-Box Optimization for Biological Sequence Design. In *International Conference on Machine Learning*, 324–334. PMLR.
- Angermueller, C.; Dohan, D.; Belanger, D.; Deshpande, R.; Murphy, K.; and Colwell, L. 2019. Model-based reinforcement learning for biological sequence design. In *International Conference on Learning Representations*.
- Aronszajn, N. 1950. Theory of reproducing kernels. *Transactions of the American mathematical society*, 68(3): 337–404.
- BakIr, G.; Hofmann, T.; Schölkopf, B.; Smola, A. J.; and Taskar, B. 2007. *Predicting structured data*. MIT press.
- Balandat, M.; Karrer, B.; Jiang, D. R.; Daulton, S.; Letham, B.; Wilson, A. G.; and Bakshy, E. 2020. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In *Advances in Neural Information Processing Systems 33*.
- Baptista, R.; and Poloczek, M. 2018. Bayesian Optimization of Combinatorial Structures. In *Proceedings of the 35th International Conference on Machine Learning*, 462–471.
- Belakaria, S.; Deshwal, A.; and Doppa, J. R. 2019. Max-value entropy search for multi-objective bayesian optimization. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Buathong, P.; Ginsbourger, D.; and Krityakierne, T. 2020. Kernels over sets of finite sets using rkhs embeddings, with application to bayesian (combinatorial) optimization. In *International Conference on Artificial Intelligence and Statistics*, 2731–2741. PMLR.
- Burkard, R. E.; Cela, E.; Pardalos, P. M.; and Pitsoulis, L. S. 1998. The quadratic assignment problem. In *Handbook of combinatorial optimization*, 1713–1809. Springer.
- Burkard, R. E.; Karisch, S. E.; and Rendl, F. 1997. QAPLIB—a quadratic assignment problem library. *Journal of Global optimization*, 10(4): 391–403.
- Che, S.; Boyer, M.; Meng, J.; Tarjan, D.; Sheaffer, J. W.; Lee, S.-H.; and Skadron, K. 2009. Rodinia: A Benchmark suite for Heterogeneous Computing. In *IEEE International Symposium on Workload Characterization (IISWC)*, 44–54. IEEE.
- Dadkhahi, H.; Shanmugam, K.; Rios, J.; Das, P.; Hoffman, S. C.; Loeffler, T. D.; and Sankaranarayanan, S. 2020. Combinatorial Black-Box Optimization with Expert Advice. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 1918–1927.
- Davis, L. 1991. *Handbook of genetic algorithms*. CumInCAD.
- Deshwal, A.; Belakaria, S.; and Doppa, J. R. 2020. Scalable Combinatorial Bayesian Optimization with Tractable Statistical models. *CoRR*, abs/2008.08177.
- Deshwal, A.; Belakaria, S.; and Doppa, J. R. 2021. Mercer Features for Efficient Combinatorial Bayesian Optimization. In *AAAI*, 7210–7218.
- Deshwal, A.; Belakaria, S.; Doppa, J. R.; and Fern, A. 2020. Optimizing discrete spaces via expensive evaluations: A learning to search framework. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 3773–3780.
- Deshwal, A.; and Doppa, J. R. 2021. Combining Latent Space and Structured Kernels for Bayesian Optimization over Combinatorial Spaces. In *Thirty-fifth Conference on Neural Information Processing Systems (NeurIPS)*.
- Deshwal, A.; Simon, C.; and Doppa, J. R. 2021. Bayesian optimization of nanoporous materials. *Molecular Systems Design and Engineering, Royal Society of Chemistry*.
- Doppa, J. R. 2021. Adaptive Experimental Design for Optimizing Combinatorial Structures. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI)*, 4940–4945.
- 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*.
- Ferreira, J. F. B.; Khoo, Y.; and Singer, A. 2018. Semidefinite programming approach for the quadratic assignment problem with a sparse graph. *Computational Optimization and Applications*, 69(3): 677–712.
- Frazier, P. I.; and Wang, J. 2016. Bayesian optimization for materials design. In *Information Science for Materials Discovery and Design*, 45–75. Springer.
- Gardner, J. R.; Pleiss, G.; Bindel, D.; Weinberger, K. Q.; and Wilson, A. G. 2018. GPyTorch: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration. In *Advances in Neural Information Processing Systems*.
- Garrido-Merchán, E. C.; and Hernández-Lobato, D. 2020. Dealing with categorical and integer-valued variables in bayesian optimization with gaussian processes. *Neurocomputing*, 380: 20–35.
- Gómez-Bombarelli, R.; Wei, J. N.; Duvenaud, D.; Hernández-Lobato, J. M.; Sánchez-Lengeling, B.; Sheberla, D.; Aguilera-Iparraguirre, J.; Hirzel, T. D.; Adams, R. P.; and Aspuru-Guzik, A. 2018. Automatic Chemical Design using a Data-driven Continuous Representation of Molecules. *ACS central science*, 4(2): 268–276.
- Greenhill, S.; Rana, S.; Gupta, S.; Vellanki, P.; and Venkatesh, S. 2020. Bayesian optimization for adaptive experimental design: A review. *IEEE Access*, 8: 13937–13948.
- Griffiths, R.-R.; and Hernández-Lobato, J. M. 2017. Constrained bayesian optimization for automatic chemical design. *arXiv preprint arXiv:1709.05501*.
- Hernández-Lobato, J. M.; Requeima, J.; Pyzer-Knapp, E. O.; and Aspuru-Guzik, A. 2017. Parallel and distributed Thompson sampling for large-scale accelerated exploration of chemical space. In *International conference on machine learning*, 1470–1479. PMLR.

- Hutter, F.; Hoos, H. H.; and Leyton-Brown, K. 2011. Sequential Model-based Optimization for General Algorithm Configuration. In *International conference on Learning and Intelligent Optimization*, 507–523.
- Jiao, Y.; and Vert, J.-P. 2015. The Kendall and Mallows kernels for permutations. In *International Conference on Machine Learning*, 1935–1944. PMLR.
- Jiao, Y.; and Vert, J.-P. 2018. The weighted kendall and high-order kernels for permutations. *arXiv preprint arXiv:1802.08526*.
- Kajino, H. 2019. Molecular hypergraph grammar with its application to molecular optimization. In *International Conference on Machine Learning*, 3183–3191. PMLR.
- Kanagawa, M.; Hennig, P.; Sejdinovic, D.; and Sriperumbudur, B. K. 2018. Gaussian processes and kernel methods: A review on connections and equivalences. *arXiv preprint arXiv:1807.02582*.
- Kendall, M. G. 1938. A new measure of rank correlation. *Biometrika*, 30(1/2): 81–93.
- Kim, J.; McCourt, M.; You, T.; Kim, S.; and Choi, S. 2021. Bayesian optimization with approximate set kernels. *Machine Learning*, 110(5): 857–879.
- Kondor, R.; and Lafferty, J. 2002. Diffusion Kernels on Graphs and other Discrete Structures. In *Proceedings of the 19th International Conference on Machine Learning*, volume 2002, 315–322.
- Korovina, K.; Xu, S.; Kandasamy, K.; Neiswanger, W.; Poczos, B.; Schneider, J.; and Xing, E. 2020. Chembo: Bayesian optimization of small organic molecules with synthesizable recommendations. In *International Conference on Artificial Intelligence and Statistics*, 3393–3403. PMLR.
- Lodhi, H.; Saunders, C.; Shawe-Taylor, J.; Cristianini, N.; and Watkins, C. 2002. Text classification using string kernels. *Journal of Machine Learning Research*, 2(Feb): 419–444.
- Lu, X.; Gonzalez, J.; Dai, Z.; and Lawrence, N. D. 2018. Structured variationally auto-encoded optimization. In *International conference on machine learning*, 3267–3275. PMLR.
- Luong, P.; Gupta, S.; Nguyen, D.; Rana, S.; and Venkatesh, S. 2019. Bayesian optimization with discrete variables. In *Australasian Joint Conference on Artificial Intelligence*, 473–484. Springer.
- Mania, H.; Ramdas, A.; Wainwright, M. J.; Jordan, M. I.; and Recht, B. 2018. On kernel methods for covariates that are rankings. *Electronic Journal of Statistics*, 12(2): 2537–2577.
- Moss, H. B.; Beck, D.; González, J.; Leslie, D. S.; and Rayson, P. 2020. BOSS: Bayesian Optimization over String Spaces. *arXiv preprint arXiv:2010.00979*.
- Murphy, K. P. 2012. *Machine learning: a probabilistic perspective*. MIT press.
- Oh, C.; Gavves, E.; and Welling, M. 2018. BOCK : Bayesian Optimization with Cylindrical Kernels. In Dy, J. G.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80 of *Proceedings of Machine Learning Research*, 3865–3874. PMLR.
- Oh, C.; Tomczak, J.; Gavves, E.; and Welling, M. 2019. Combinatorial Bayesian Optimization using the Graph Cartesian Product. In *Advances in Neural Information Processing Systems*, 2910–2920.
- Rasmussen, C. E.; and Williams, C. K. I. 2006. *Gaussian processes for machine learning*. Adaptive computation and machine learning. MIT Press.
- Reinelt, G. 1995. Tsplib95. *Interdisziplinäres Zentrum für Wissenschaftliches Rechnen (IWR), Heidelberg*, 338: 1–16.
- Russo, D.; and Van Roy, B. 2014. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 39(4): 1221–1243.
- Sahni, S.; and Gonzalez, T. 1976. P-complete approximation problems. *Journal of the ACM (JACM)*, 23(3): 555–565.
- Shahriari, B.; Swersky, K.; Wang, Z.; Adams, R. P.; and De Freitas, N. 2016. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1): 148–175.
- 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*.
- Swersky, K.; Rubanova, Y.; Dohan, D.; and Murphy, K. 2020. Amortized bayesian optimization over discrete spaces. In *Conference on Uncertainty in Artificial Intelligence*, 769–778. PMLR.
- Varga, R. S. 2010. *Geršgorin and his circles*, volume 36. Springer Science & Business Media.
- Yang, K. K.; Wu, Z.; and Arnold, F. H. 2019. Machine-learning-guided directed evolution for protein engineering. *Nature methods*, 16(8): 687–694.
- Zhao, Q.; Karisch, S. E.; Rendl, F.; and Wolkowicz, H. 1998. Semidefinite programming relaxations for the quadratic assignment problem. *Journal of Combinatorial Optimization*, 2(1): 71–109.
- Zhou, Z.; Belakaria, S.; Deshwal, A.; Hong, W.; Doppa, J. R.; Pande, P. P.; and Heo, D. 2020. Design of Multi-Output Switched-Capacitor Voltage Regulator via Machine Learning. In *DATE*.

8 Additional Experiments

optimization for finding the best structure.

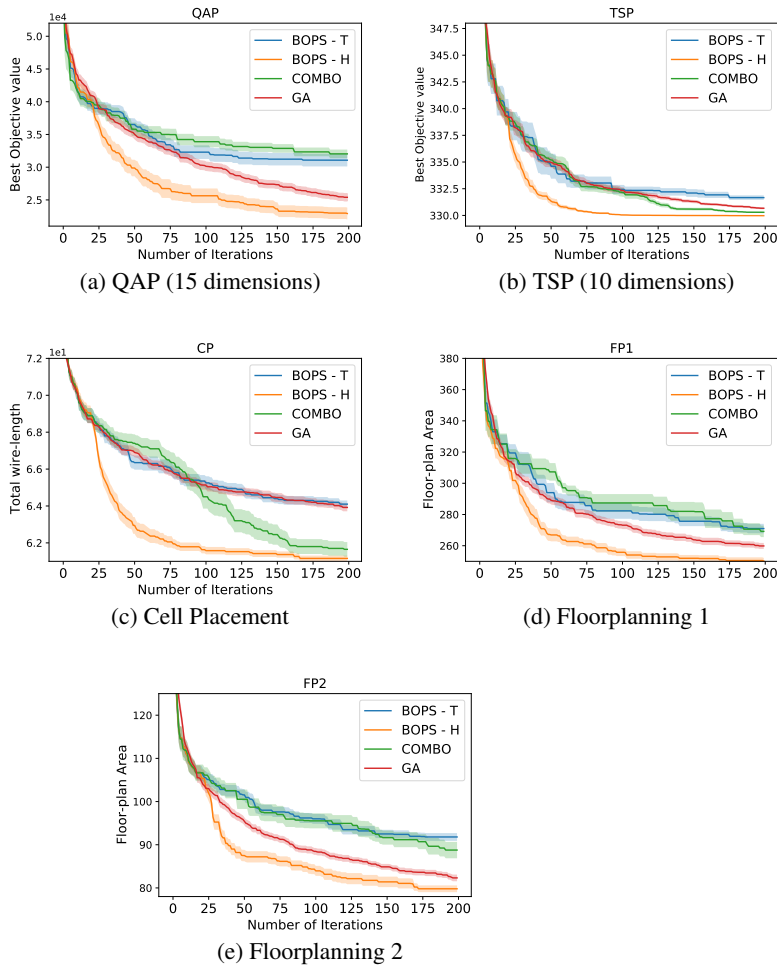


Figure 3: Results comparing BOPS-T, BOPS-H, and COMBO and Genetic Algorithm (GA) (best objective function value vs. number of BO iterations) on both synthetic and real-world benchmarks:

Genetic algorithms (Davis 1991) are a commonly employed technique for optimizing discrete structures. They have also been used for optimizing acquisition functions (Moss et al. 2020) over discrete spaces in Bayesian optimization. We present additional results comparing our proposed algorithms with genetic algorithms in Figure 3. We used a genetic algorithm with permutation based operators as implemented in Platypus <https://github.com/Project-Platypus/Platypus> library. We tried few different choices for the two hyper-parameters (population size and offspring size) and picked population size to be 20 and offspring size to be 10. Since we are in a small data setting, large values for these hyper-parameters will not be practically useful. As seen in the figure, BOPS-H still performs the best outperforming even genetic algorithms demonstrating the importance incorporating a model-guided search like Bayesian