Design-Bench: Benchmarks for Data-Driven Offline Model-Based Optimization

Brandon Trabucco, Aviral Kumar, Young Geng, Sergey Levine
Department of Computer Science
UC Berkeley
Berkeley, CA 94720, USA
{btrabucco,aviralk,young.geng,sergey.levine}@berkeley.edu

Abstract

Black-box model-based optimization (MBO) problems, where the goal is to find a design input that maximizes an unknown objective function, are ubiquitous in a wide range of domains, such as the design of drugs, aircraft, and robot morphology. Typically, such problems are solved by actively querying the black-box objective on design proposals and using the resulting feedback to improve the proposed designs. However, when the true objective function is expensive or dangerous to evaluate in the real world, we might instead prefer a method that can optimize this function using only previously collected data, for example from a set of previously conducted experiments. This data-driven offline MBO setting presents a number of unique challenges, but a number of recent works have demonstrated that viable offline MBO methods can be developed even for high-dimensional problems, using high-capacity deep neural network function approximators. Unfortunately, the lack of standardized evaluation tasks in this emerging new field has made tracking progress and comparing recent methods difficult. To address this problem, we present Design-Bench, a benchmark suite of offline MBO tasks with a unified evaluation protocol and reference implementations of recent methods. Our benchmark suite includes diverse and realistic tasks derived from real-world problems in biology, materials science, and robotics that present distinct challenges for offline MBO methods. Our benchmarks, together with the reference implementations, are available at sites.google.com/view/design-bench. We hope that our benchmark can serve as a meaningful metric for the progress of offline MBO methods and guide future algorithmic development.

1 Introduction

Automatically synthesizing designs that maximize a desired objective function is one of the most important problems in many scientific and engineering domains. From protein design in molecular biology [1] to superconducting material discovery in physics [2], researchers have made significant progress in applying machine learning methods to such optimization problems over structured design spaces. Commonly, the exact form of the objective function is unknown, and the objective values for a novel design can only be evaluated by running either computer simulations or physical experiments in the real world. The process of optimizing an unknown function is known as black-box optimization, and is typically solved in an **online** iterative manner, where in each iteration the solver proposes new designs and query the objective function for feedback in order to propose better design in the next iteration [3]. In many domains however, the evaluation of the objective function is prohibitively expensive, because it requires manually conducting experiments in the real world. In this setting, one cannot simply query the true objective function to gradually improve the design. Instead, a collection of past records of designs and their corresponding objective values might be available, and

Machine Learning for Structural Biology Workshop, NeurIPS 2020, Vancouver, Canada.

therefore the optimization method must leverage the available data to synthesize the most optimal design possible. This is the setting of data-driven **offline model-based optimization**.

Although online black-box optimization has been studied extensively, the offline MBO problem has received comparatively less attention, and only a small number of recent works study offline MBO in the setting with high-dimensional design spaces, where they utilize deep learning techniques [4, 5, 6]. This is partly due to the fact that methods for online design optimization cannot be easily applied in the offline MBO setting. However, even with only a few existing methods, it is still hard to compare and track the progress in this field, as these methods are proposed and evaluated on different tasks with distinct evaluation protocols. To the best of our knowledge, there is no commonly adopted set of benchmarks for offline MBO. To address this problem, in this paper we introduce a suite of offline MBO benchmarks with standardized evaluation protocols. We include a realistic and diverse set of tasks that spans a wide range of application domains, from synthetic biology to robotics. The realism and diversity of the tasks is essential for the evaluation of offline data-driven model-based optimization methods, as it measures the generality of the methods being evaluated across multiple domains and verifies that they are not overfitting to a single task. Our benchmark tasks incorporate a variety of challenging factors, such as high dimensional input spaces and sensitive discontinuous objective functions, which help better identify the strengths and weaknesses of MBO methods.

Along with the benchmark suite, we also present reference implementations of a number of existing offline MBO methods and baselines. We systematically evaluate them on all of the proposed benchmark tasks and report our findings. A surprising discovery from our findings is that with proper data normalization, the simple baseline method of learning an objective value predictor and performing gradient ascent on its input outperforms several prior MBO methods in our benchmark. We hope that our work can provide insight into the current progress of offline MBO methods and can also serve as a meaningful metric to galvanize research in this area.

2 Offline Model-Based Optimization Problem Statement

The goal in offline model-based optimization is to optimize an unknown (possibly stochastic) objective function $f(\mathbf{x})$, provided access to a static dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}$ of designs $\mathbf{x}_i$ and a corresponding measurement of the objective value y_i. Similar to batch Bayesian optimization (BBO) [7], each algorithm $\mathfrak{A}$ is allowed to consume the dataset $\mathcal{D}$, and is required to produce a set of K candidate designs $\mathfrak{A}(\mathcal{D}, K) = \{\mathbf{x}_i^* : i \in \{1...K\}\}$. These K candidates are evaluated under the ground truth objective function $f(\mathbf{x})$, and the best performing design is reported as the final performance value. Abstractly, the objective for offline MBO is:

$$\arg\max_{\mathfrak{A}} \left[\mathbb{P}(\{f(\mathbf{x}^*) : \mathbf{x}^* \in \mathfrak{A}(\mathcal{D}, K)\}, N)\right], \tag{1}$$

where $\mathbb{P}$ denotes the percentile function. Intuitively, this formulation ranks an offline MBO algorithm using the N^{th} percentile objective value obtained by it given a fixed evaluation budget of $K-$samples. Common choices of N are 100, which represent the max objective value, and 50 which represents the median objective value among the candidates.

What makes offline MBO especially challenging? The offline nature of the problem requires that the algorithm $\mathfrak{A}$ not be tuned by peeking into the ground truth objective f, and this makes the offline MBO problem much more difficult than the online design optimization problem. One simple idea to tackle this problem is to learn an objective proxy using the dataset, and then convert this offline MBO problem into an online problem by treating the learned objective proxy as the true objective. However, this idea may not work well, due to the intrinsic out-of-distribution nature of optimal designs. First of all, in a number of practical MBO problems such as optimization over proteins or robot morphologies, the designs with highest objective values in the dataset already lie on the tail of the dataset distribution, since they are better than most other designs. In order to improve upon the best designs, an optimization method needs to produce designs that are even further away from the dataset distribution. For such out-of-distribution designs, is would be impossible to guarantee that the learned objective proxy is accurate, and hence any powerful optimization method would easily "exploit" the learned objective proxy and produce falsely promising designs that are drastically overestimated by the learned objective proxy. This conflict between the out-of-distribution nature of optimization and the in-distribution requirement of any learned model is indeed the core challenge of offline MBO. This challenge is often exacerbated in real-world problems by the high dimensionality of the design space and the sparsity of the available data, as we will show in our benchmark. A good

Task	Dataset Size	Design Dimensions	Design Space	Discrete Categories
GFP-v0	5000	238	Discrete	20
MoleculeActivity-v0	4216	1024	Discrete	2
Superconductor-v0	16953	81	Continuous	N/A
HopperController-v0	3200	5126	Continuous	N/A
AntMorphology-v0	12300	60	Continuous	N/A
DKittyMorphology-v0	9546	56	Continuous	N/A

Table 1: **Overview of the tasks in our benchmark suite.** Design-Bench includes a variety of tasks from different domains with both discrete and continuous design spaces and 3 high-dimensional tasks with > 200 design dimensions, making it suitable for benchmarking offline MBO methods.

offline MBO method needs to carefully balance the two sides of the conflict, producing optimized designs that are not too far from the data distribution.

3 Related Work

While the extensive prior work on MBO, particularly Bayesian optimization, has made tremendous progress both in terms of enabling intelligent active selection of query points [8, 9, 10, 11, 12] and scaling up the dimensionality of Bayesian optimization methods [13, 14, 15], a significant number of real-world problems may be approached more naturally as offline MBO problems, given the availability of several datasets which contain designs annotated with their corresponding score values. Offline MBO presents a new set up challenges, as we discuss below, and we believe that these challenges require a new set of benchmarks that emphasize offline data and high-dimensional design spaces.

Researchers working on either online active design optimization or offline model-based optimization have collected various datasets of designs that can be used to build tasks for offline MBO. sarkisyan2016GFP analyze the fluorescence of GFP proteins under blue and ultraviolet light, and brookes2019conditioning use this dataset for optimization to find the protein with the highest fluorescence value. ChEMBL [16] provides a dataset for drug discovery, where molecule activities are measured against a target assay. hamidieh2018superconductor analyze the critical temperatures for superconductors and provide a dataset to search for room-temperature superconductors with potential in the construction of quantum computers. Some of these datasets have already been employed in the study of offline MBO methods [5, 4, 6]. However, these studies use different set of tasks and their evaluation protocols are highly field-specific, making it difficult to use for general algorithm development. In our benchmark, we incorporate modified variants of some of these datasets along with our own tasks and provide a standardized evaluation protocol. We hope that these tasks can represent realistic MBO problems across a wide range of domains and that the standardized evaluation protocol can facilitate development of new and more powerful offline MBO algorithms.

Recently, researchers have proposed several methods specifically for the completely offline MBO problem. To make up for the fact that the true objective function is not available, these methods often make use of some type of learned objective function as a proxy for the real one and a generative model [17, 18] to capture the distribution of valid designs, especially for high dimensional design spaces. [5] learns a objective value conditioned generative model to synthesize the designs with the desired objective value. [4, 6] employ a variational autoencoder [17] to capture the distribution of designs in the dataset and a learned neural network objective function model. In our benchmark, we provide open-source reference implementations for these methods and systematically evaluate them in all our tasks.

4 Design-Bench Benchmark Tasks

In this section, we describe the set of tasks included in our benchmark. We first provide an overview of the tasks in Table 1. Each task in our benchmark suite comes with a dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}$, along with an oracle objective function f which is to be utilized for for test time evaluation only. An offline MBO algorithm should not query the oracle during training time, even for hyperparameter tuning. While in some of the tasks in our benchmark, such as tasks pertaining to robotics – HopperController-v0, DKittyMorphology-v0 and AntMorphology-v0, the oracle functions are evaluated by running computer simulations to obtain the true objective values, in the other tasks, the true objective values can only be obtained by conducting expensive physical experiments. While the eventual aim of offline MBO is to make it possible to optimize designs in precisely such settings where querying the groundtruth objective is challenging, requiring real physical experiments for evaluation makes the design and benchmarking of new algorithms difficult and time consuming. Therefore, to facilitate evaluation, following evaluation methodology in prior work [4, 6], and use models built by experts as our oracle function. Such models are typically *also* learned models, but with representations that are hand-designed, built-in domain-specific inductive biases and typically trained on much more data than is made available for solving the offline MBO problem, which increases the chance that this proxy "true function" can answer queries outside of the MBO training distribution. While this approach to evaluation diminishes the realism of our benchmark, since these proxy "true functions" may not always be accurate, we believe that this trade off is worth it to make benchmarking practical. The main purpose of our benchmark is to facilitate the evaluation and development of offline MBO algorithms, and we believe that it is important to include tasks in domains where the true objective values can only be obtained via physical experiments, which make up a large portion of the real-world MBO problems. For each task we discuss in this section, we provide a detailed description of the data collection strategy and the data pre-processing strategy in Appendix A.

GFP-v0: protein fluorescence maximization. The goal of this task is to design a derivative protein from the *Aequorea victoria* green fluorescent protein (GFP) that has maximum fluorescence, using a real-world dataset mapping proteins to fluorescence collected by [19]. While we cannot exactly evaluate any novel protein, we employ an expert Gaussian process regression model with a protein-specific kernel built by [1] as the oracle function, following the convention of [4, 6]. This expert model is built on a larger dataset than the one we use in this benchmark, making it relatively accurate for proteins not in our dataset. The design space for GFP is discrete, consisting of a sequence of 238 categorical variables that can take one of 20 options, which corresponds to the choice of 238 amino acids that make up the green fluorescent protein from 20 types of amino acid.

MoleculeActivity-v0: drug activity maximization. This task is taken from the domain of drug discovery, where the goal is to design the substructure of a molecule that exhibits high activity when tested against a target assay. The dataset we provide with this task was originally collected by ChEMBL [16]. Again like for the GFP task, the true molecule activity against a target assay can only be evaluated with physical experiments. Therefore, we leverage an oracle used recently for meta-learning [20] and adopt the expert random forest regression model built by [21] as the oracle function. From the ChEMBL data, we pick a single target assay to form the dataset of this task, resulting in 4216 molecules in total. The design space for this task is a sequence of 1024 binary categorical variables that correspond to the Morgan radius 2 substructure fingerprints, making this task a high-dimensional task.

Superconductor-v0: critical temperature maximization for superconductor materials. The Superconductor-v0 task is taken from the domain of materials science, where the goal is to design a superconducting material that has a high critical temperature. Developing superconducting materials has immense practical value, especially for room-temperature superconductors, which could greatly facilitate quantum computing. We adapt a real-world dataset used by materials scientists [2]. The dataset contains 21263 superconductors annotated with critical temperatures. Prior work has employed this dataset for the study of offline MBO methods [4], and we follow their convention of using an expert random forest model, detailed in [6], for our oracle function. The design space for Superconductor-v0 is an 81-channel vector of continuous variables, which describe properties of the material such as its atomic radius.

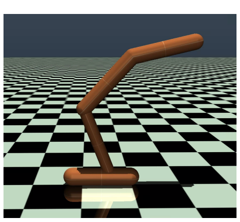

HopperController-v0: robot neural network controller optimization. HopperController-v0 is a task created by us in the domain of robotics, where the goal is to optimize weights of a neural network policy that maximizes expected return on the Hopper-v2 task in OpenAI Gym [22]. While this goal might look similar to that of reinforcement learning, it is important to note that the formulation is entirely different. Unlike reinforcement learning, we don't have access to any form of trajectory data, neither by actively sampling from the environment nor from a pre-collected dataset. Instead, we only have a dataset of weights of a neural network controller and the corresponding return values. Therefore, reinforcement learning methods cannot be applied to solve this problem. Unlike tasks introduced before, we evaluate the true objective value of any design during dataset generation by running 1000 steps of simulation in the MuJoCo simulator. This horizon length is chosen so that the learning problem is long-horizon, and difficult. The design space of this task is high-dimensional with 5126 continuous variables corresponding to the flattened weights of a neural network controller. The dataset is collected by training a PPO reinforcement learning agent [23] and recording the agent's weights every 10,000 samples.

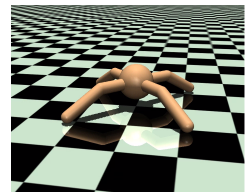

AntMorphology-v0 and DKittyMorphology-v0: robot morphology optimization. We created these two tasks to optimize the morphological structure of two simulated robots: Ant from OpenAI Gym [22] and D'Kitty from ROBEL [24]. For AntMorphology, the goal is to optimize the morphology of an ant shaped robot, to run as fast as possible, with a pre-trained neural network controller. For DKittyMorphology, the goal is to optimize the morphology of D'Kitty, a quadrupedal robot, such that a pre-trained neural network controller can navigate the robot to a fixed location. In this fashion, the goal for both tasks is to *recover* a morphology with which the pre-trained controller is compatible. The variable morphology parameters of both robots include size, orientation, and location of the limbs, giving us 60 continuous values in total for Ant and 56 for D'Kitty. To evaluate the ground truth value for a given design, we run robotic simulation in MuJoCo [25] for 100 time steps, averaging 16 independent trials. These parameters are chosen to reduce stochasticity, and allow the simulator to run in a minimal amount of time.

5 Task Properties, Challenges, and Considerations

The primary goal of our proposed benchmark is to provide a general test bench for developing, evaluating, and comparing algorithms for offline MBO. While in principle any online active design optimization problem can be formulated into an offline MBO problem by collecting a dataset of designs and corresponding objective measurements, it is important to pick a subset of tasks that represent the challenges of real-world MBO problems in order to convincingly evaluate algorithms and obtain insights about algorithm behavior. Therefore, several factors must be considered when choosing the tasks, which we discuss next.

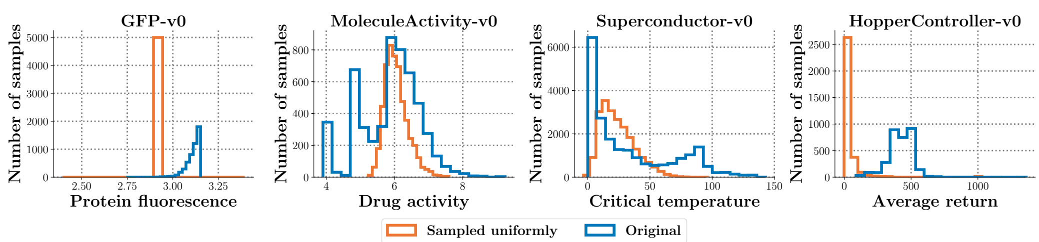

Figure 1: **Histogram (frequency distribution) of objective values in the dataset compared to a uniform re-sampling of the dataset** from the design space. In every case, re-sampling skews the distribution of scores to the left, suggesting that there exists a thin manifold of valid designs in the high-dimensional design space, and most of the volume in this space is occupied by low-scoring designs.

Diversity and Realism. First of all, the tasks need to be diverse and realistic in order to prevent offline MBO algorithms from overfitting to a particular problem domain and to expect that offline MBO methods performing well on this benchmark suite would also perform well on other real-world offline MBO problems. Design-Bench consists of tasks diverse in many respects: it includes both tasks with *discrete* and with *continuous* design spaces, which have distinct implications for offline MBO algorithms. Continuous design spaces, equipped with metric space and ordering structures, could make the problem easier to solve than discrete design spaces. On the other hand, discrete design spaces are finite and therefore the dataset coverage might be better than some continuous tasks. A strong offline MBO algorithm needs to be able to handle both cases. Moreover, our tasks have varying dimensionality, ranging from 56 to 5126 dimensions. And finally, our tasks span multiple realistic domains ranging from biology and physical science to robotics.

High-dimensional design space. In many real-world offline MBO problems such as drug discovery [16], the design space is *high-dimensional* and the valid designs lie on a *thin manifold* in this high-dimensional space. Due to the curse of dimensionality, the available design data is often sparsely distributed in the design space. This property poses a unique challenge for many MBO algorithms: to

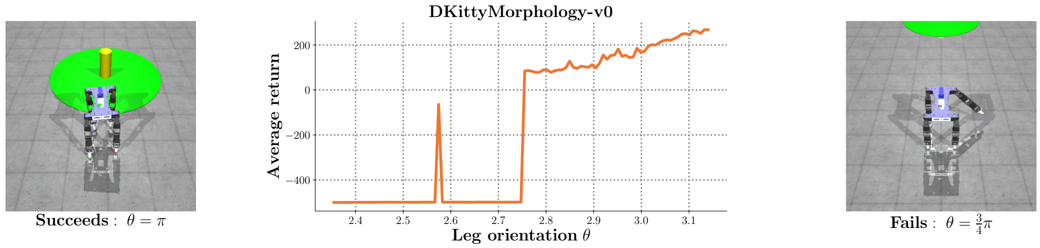

Figure 2: **Highly sensitive landscape of the groundtruth objective function in DKittyMorphology-v0.** A small change in a single dimension of the design-space, for instance, by changing the orientation θ (x-axis) of the base of the agent's front right leg, the performance value (y-axis) of the agent is critically impacted. The agent's design is the original D'Kitty design and is held constant while varying θ uniformly from $\frac{3}{4}\pi$ to π.

be effective on such problem domains, MBO methods need to strongly restrict optimization to a thin manifold of the design space to be able to produce valid designs. Prior work [5] has noted that this can be extremely hard in several cases. In our benchmark, we include the GFP-v0, MoleculeActivity-v0 and HopperController-v0 tasks with high-dimensional design spaces to capture this challenge. To intuitively understand this challenge, we performed a study on the HopperController-v0 task in Figure 1, where we sampled 3200 designs uniformly at random from the design space and plotted a histogram of their objective values against those in the dataset we provide, that only consists of valid designs. Observe the clear discrepancy in the objective values, where randomly sampled designs attain objective values much lower than the dataset. This indicates that valid designs only lie on a thin manifold in the design space and therefore we are very unlikely to hit a valid design by uniform random sampling.

Highly sensitive objective function. Another important challenge that should be taken into consideration is the *high sensitivity* of objective functions, where closeness of two designs in design space need not correspond to closeness in their objective values, which may differ drastically. This challenge is naturally present in practical problems like protein synthesis [1], where the change of a single amino acid could significantly alter the property of the protein. The DKittyMorphology-v0 and AntMorphology-v0 tasks in our benchmark suite are also particularly challenging in this direction. To visualize the high sensitivity of the objective function, we plot a one dimensional slice of the objective function around a single sample in our dataset in Figure 2. We see that with other variables kept the same, slightly altering one variable can significantly reduce the objective value, making it hard for offline MBO methods to produce the optimal design.

6 Algorithm Implementations In Design-Bench

To provide a baseline for comparisons in future work, we ran a number of recently proposed MBO algorithms on each of our tasks. Since the dimensionality of our tasks ranges from 56 to 5126, we chose methods that have been shown in prior work to handle *both* the case of offline training data (i.e., no active interaction) and high-dimensional inputs. To that end, we include MINs [5], CbAS [4], and autofocusing CbAS [6] in our comparisons, along with a baseline naïve "gradient ascent" method that approximates the true function $f(\mathbf{x})$ with a deep neural network and then performs gradient ascent on the output of this model. In this section, we briefly discuss these algorithms, before performing a comparative evaluation in the next section.

Gradient ascent baseline. We evaluate a simple baseline that learns a model of the objective function, $\hat{f}(\mathbf{x})$ and optimizes $\mathbf{x}$ against this learned model via gradient ascent. Formally, the optimal solution $\mathbf{x}^*$ generated by this method can be computed as a fixed point of the following update: $\mathbf{x}_{t+1} \leftarrow \mathbf{x}_t + \alpha \nabla_{\mathbf{x}} \hat{f}(\mathbf{x})|_{\mathbf{x}=\mathbf{x}_t}$. In practice we perform T gradient steps, and report $\mathbf{x}_T$ as the final solution. Such methods are notoriously susceptible to falling off-the-manifold of valid inputs since nothing constrains the resulting $\mathbf{x}_T$ to be on the manifold of valid-designs.

Conditioning by adaptive sampling (CbAS). CbAS learns a density model in the space of design inputs, $p_0(\mathbf{x})$ that approximates the data distribution and gradually adapts it towards the optimized solution $\mathbf{x}^*$. In a particular iteration t, CbAS alternates between **(1)** training a variational auto-encoder (VAE) [17] on a set of samples generated from the previous model $\mathcal{D}_t = \{\mathbf{x}_i\}_{i=1}^m; \mathbf{x}_i \sim p_{t-1}(\cdot)$ using a weighted version of the standard ELBO objective that puts a higher likelihood on *estimated* high-scoring designs (i.e., designs that have a score beyond a certain pre-specified threshold), and **(2)** generating new design samples from the autoencoder to serve as $\mathcal{D}_{t+1} = \{\mathbf{x}_i | \mathbf{x}_i \sim p_t(\cdot)\}$. In order to

Task Identifier	$\mathcal{D}$ (best)	MINs	CbAS	Autofocus	Gradient ascent
GFP-v0	3.152	3.315 ± 0.033	3.408 ± 0.029	3.365 ± 0.023	2.894 ± 0.001
MoleculeActivity-v0	6.558	6.508 ± 0.236	6.301 ± 0.131	6.345 ± 0.141	6.636 ± 0.066
Superconductor-v0	73.90	80.23 ± 10.67	72.17 ± 8.652	77.07 ± 11.11	89.64 ± 9.201
HopperController-v0	1361.6	746.1 ± 636.8	547.1 ± 423.9	443.8 ± 142.9	1050.8 ± 284.5
AntMorphology-v0	108.5	388.5 ± 9.085	393.0 ± 3.750	386.9 ± 10.58	399.9 ± 4.941
DKittyMorphology-v0	215.9	352.9 ± 38.65	369.1 ± 60.65	376.3 ± 47.47	390.7 ± 49.24

Table 2: **100th percentile** evaluations for baselines on every task. Results are averaged over 16 trials, and the $\pm$ indicates the standard deviation of the reported performance.

Task Identifier	$\mathcal{D}(50^{\text{th}})$	MINs	CbAS	Autofocus	Gradient ascent
GFP-v0	3.121	3.135 ± 0.019	3.269 ± 0.018	3.216 ± 0.029	2.894 ± 0.000
MoleculeActivity-v0	5.876	5.806 ± 0.078	5.742 ± 0.123	5.759 ± 0.158	6.401 ± 0.186
Superconductor-v0	12.00	37.32 ± 10.50	32.21 ± 7.255	31.57 ± 7.457	54.06 ± 5.060
HopperController-v0	434.5	520.4 ± 301.5	132.5 ± 23.88	116.4 ± 18.66	185.0 ± 72.88
AntMorphology-v0	37.44	184.8 ± 29.52	267.3 ± 16.55	176.7 ± 59.94	318.0 ± 12.05
DKittyMorphology-v0	190.2	211.6 ± 13.67	203.2 ± 3.580	199.3 ± 8.909	255.3 ± 6.379

Table 3: **50th percentile** evaluations for baselines on every task.

estimate the objective values for design samples from the learned density model $p_t(\mathbf{x})$, CbAS utilizes separately trained models of the objective function, $\hat{f}(\mathbf{x})$, that are trained via supervised regression to map an input $\mathbf{x}_i$ to its objective value y_i. This training process, at a given iteration t, can formally be described as:

$$p_{t+1}(\mathbf{x}) := \arg\min_{p} \frac{1}{m} \sum_{i=1}^{m} \frac{p_0(\mathbf{x}_i)}{p_t(\mathbf{x}_i)} P(\hat{f}(\mathbf{x}_i) \geq \tau) \log p_t(\mathbf{x}_i), \quad \text{where} \;\; \{\mathbf{x}_i\}_{i=1}^{m} \sim p_t(\cdot). \tag{2}$$

Autofocusing CbAS. Since CbAS uses a learned model of the objective function $\hat{f}(\mathbf{x})$ as a proxy for the ground-truth function to iteratively adapt the generative model $p(\mathbf{x})$ towards the optimized design, the function$\hat{f}(\mathbf{x})$ will inevitably be required to make predictions under shifting design distributions $p_t(\mathbf{x})$. Since these learned predictions are used as ground-truth values, any inaccuracy in the learned function can adversely affect the optimization procedure. Autofocused CbAS aims to correct for this distribution shift by re-training $\hat{f}(\mathbf{x})$ (now denoted $\hat{f}_t(\mathbf{x})$) under the design distribution given by the current model, $p_t(\mathbf{x})$ via importance sampling, which is then used by CbAS.

$$\hat{f}_{t+1} := \arg\min_{\hat{f}} \;\; \frac{1}{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{D}|} \frac{p_t(\mathbf{x}_i)}{p_0(\mathbf{x}_i)} \cdot \left(\hat{f}(\mathbf{x}_i) - y_i\right)^2,$$

Model inversion networks (MINs). MINs learn an inverse map from the objective value to a design, $\hat{f}^{-1} : \mathcal{Y} \rightarrow \mathcal{X}$ by using objective-conditioned inverse maps, search for optimal y values during optimization and finally query the learned inverse map to produce the corresponding optimal design. These methods minimize a divergence measure $\mathcal{L}_p(\mathcal{D}) := \mathbb{E}_{y \sim p_{\mathcal{D}}(y)} \left[D(p_{\mathcal{D}}(\mathbf{x}|y), \hat{f}^{-1}(\mathbf{x}|y))\right]$, to train such an inverse map. During optimization, MINs obtain the optimal y-value that is used to query the inverse map, and obtains the optimized design by sampling form the inverse map.

7 Benchmarking Prior Methods

In this section, we provide a comparison of prior algorithms discussed in Section 6 on our proposed tasks. For purposes of standardization, easy benchmarking, and future algorithm development, we present results for all Design-Bench tasks in Table 2 and 3. As discussed in Section 2, we provide each algorithm with a dataset, and allow the method to produce $K = 128$ optimized design candidates. These $K = 128$ candidates are then evaluated with the oracle function, and we report the 100^{th} percentile and the 50^{th} percentile performance among them averaged over 16 independent trials, following the convention of prior offline MBO works [6].

Algorithm setup and hyperparameter tuning. Since our goal is to generate high-performing solutions without *any* knowledge of the groundtruth, oracle function, any form of hyperparameter tuning on the parameters of the learned model should respect the evaluation boundary. We provide a recommended method for tuning each algorithm described in Section 6 that is fully offline. To briefly summarize, for **CbAS**, tuning amounts to finding a stable configuration for a VAE, such that samples from the prior distribution map to on-manifold designs after reconstruction. We empirically found that a β-VAE was essential for stability of CbAS—and high values of $\beta > 1$ are especially important for modelling high-dimensional spaces like that of HopperController-v0. As a general task-agnostic principle for selecting β, we choose the smallest β such that the VAE's latent space does not collapse during importance sampling. Collapsing latent-spaces seem to coincide with diverging importance sampling, and the VAE's reconstructions collapsing to a single mode. For **MINs**, tuning amounts to fitting a good generative model. We observe that MINs is particularly sensitive to the scale of y_i when conditioning, which we resolve by normalizing the objective values. We implement MINs using WGAN-GP, and find that similar hyperparameters work well-across domains. For **Gradient Ascent**, while prior work has generally obtained extremely poor performance for naïve gradient-ascent based optimization procedures on top of learned models of the objective function, we find that by normalizing the designs $\mathbf{x}$ and objective values y to have unit Gaussian statistics, and by multiplying the learning rate $\alpha \leftarrow \alpha\sqrt{d}$ where d is the dimension of the design space, a naïve gradient ascent based procedure can perform reasonably well on all tasks, without any task-specific tuning. For discrete tasks, only the objective values are normalized, and optimization is performed over log-probabilities of designs. We then uniformly evaluate samples obtained by running 200 steps of gradient ascent starting from the highest scoring 128 samples in each dataset.

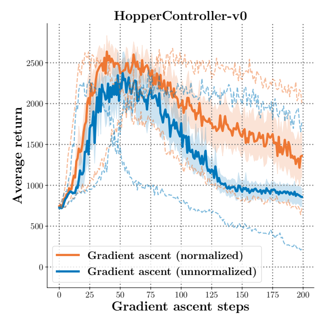

Results and discussion. The results for all tasks are provided in Table 2 and 3. There are several interesting takeaways from these results. First, these results indicate that there is no clear winner between the three prior methods MINs, CbAS and Autofocused CbAS, provided they are all trained offline with no access to groundtruth evaluation for any form of hyperparameter tuning. Second, somewhat surprisingly, the naïve gradient ascent baseline outperforms most of these highly sophisticated MBO methods based on generative models in **5 out of 6** tasks (Table 2), especially on high-dimensional tasks (e.g., HopperController-v0). This result suggests that it might be difficult for the generative models used in prior method to capture the distribution of the dataset. We conducted an ablation study to determine the reasons behind the surprisingly good performance of the naïve baseline. We found that the identical gradient-ascent baseline performed a factor **1.5x** worse on HopperController-v0, when optimizing in the space of unnormalized designs and objective values, indicating that normalization is key in obtaining good performance with a naïve gradient ascent baseline.

Finally, we remark that the performance numbers for certain methods reported in Table 2 differ from the performance reported by prior works on certain tasks This difference stems from the standardization procedure employed in dataset generation (which we discuss in Appendix A), and the usage of uniform hyperparameters to ensure task-agnostic hyperparameter selection.

8 Conclusion

While online model-based optimization methods have received a lot of attention, especially in the Bayesian optimization community, offline model-based optimization is increasingly gaining popularity as it carries the promise to be able to convert existing databases of designs into powerful optimizers, without the need for any expensive online interaction. However, due to the lack of standardized benchmarks and evaluation protocols, it has been difficult to accurately track the progress in this field. To address this problem, we introduce Design-Bench, a benchmark suite of offline MBO tasks that covers a wide variety of domains, and both continuous and discrete, low and high dimensional design spaces. We provide a comprehensive evaluation of existing methods under identical assumptions. The somewhat surprising efficacy of simple baselines such as naïve gradient ascent suggests the need for careful tuning and standardization of methods in this area. We hope that our benchmark will be adopted as the standard metric in evaluating offline MBO algorithms and provide meaningful insight in future algorithmic development.

References

[1] W.-J. Shen, H.-S. Wong, Q.-W. Xiao, X. Guo, and S. Smale, "Introduction to the peptide binding problem of computational immunology: New results," *Foundations of Computational Mathematics*, vol. 14, pp. 951–984, Oct 2014.

[2] K. Hamidieh, "A data-driven statistical model for predicting the critical temperature of a superconductor," *Computational Materials Science*, vol. 154, pp. 346 – 354, 2018.

[3] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*, vol. 2. MIT press Cambridge, MA, 2006.

[4] D. H. Brookes, H. Park, and J. Listgarten, "Conditioning by adaptive sampling for robust design," *arXiv preprint arXiv:1901.10060*, 2019.

[5] A. Kumar and S. Levine, "Model inversion networks for model-based optimization," *arXiv preprint arXiv:1912.13464*, 2019.

[6] C. Fannjiang and J. Listgarten, "Autofocused oracles for model-based design," *arXiv preprint arXiv:2006.08052*, 2020.

[7] J. González, Z. Dai, P. Hennig, and N. Lawrence, "Batch bayesian optimization via local penalization," in *Artificial intelligence and statistics*, pp. 648–657, 2016.

[8] T. G. Kolda, R. M. Lewis, and V. Torczon, "Optimization by direct search: New perspectives on some classical and modern methods," *SIAM review*, vol. 45, no. 3, pp. 385–482, 2003.

[9] D. J. Lizotte, *Practical bayesian optimization*. University of Alberta, 2008.

[10] R. Y. Rubinstein and D. P. Kroese, *The cross-entropy method: a unified approach to combinatorial optimization, Monte-Carlo simulation and machine learning*. Springer Science & Business Media, 2013.

[11] M. J. Powell, "Direct search algorithms for optimization calculations," *Acta numerica*, pp. 287–336, 1998.

[12] D. Whitley, "A genetic algorithm tutorial," *Statistics and computing*, vol. 4, no. 2, pp. 65–85, 1994.

[13] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," in *Advances in neural information processing systems*, pp. 2951–2959, 2012.

[14] J. Snoek, O. Rippel, K. Swersky, R. Kiros, N. Satish, N. Sundaram, M. Patwary, M. Prabhat, and R. Adams, "Scalable bayesian optimization using deep neural networks," in *International conference on machine learning*, pp. 2171–2180, 2015.

[15] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas, "Taking the human out of the loop: A review of bayesian optimization," *Proceedings of the IEEE*, vol. 104, no. 1, pp. 148–175, 2015.

[16] A. Gaulton, L. J. Bellis, A. P. Bento, J. Chambers, M. Davies, A. Hersey, Y. Light, S. McGlinchey, D. Michalovich, B. Al-Lazikani, and J. P. Overington, "Chembl: a large-scale bioactivity database for drug discovery," *Nucleic acids research*, vol. 40, pp. D1100–D1107, Jan 2012. 21948594[pmid].

[17] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.

[18] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, pp. 2672–2680, 2014.

[19] K. S. Sarkisyan, D. A. Bolotin, M. V. Meer, D. R. Usmanova, A. S. Mishin, G. V. Sharonov, D. N. Ivankov, N. G. Bozhanova, M. S. Baranov, O. Soylemez, N. S. Bogatyreva, P. K. Vlasov, E. S. Egorov, M. D. Logacheva, A. S. Kondrashov, D. M. Chudakov, E. V. Putintseva, I. Z. Mamedov, D. S. Tawfik, K. A. Lukyanov, and F. A. Kondrashov, "Local fitness landscape of the green fluorescent protein," *Nature*, vol. 533, pp. 397–401, May 2016.

[20] H. Yao, L. Huang, Y. Wei, L. Tian, J. Huang, and Z. Li, "Don't overlook the support set: Towards improving generalization in meta-learning," *CoRR*, vol. abs/2007.13040, 2020.

[21] E. J. Martin, V. R. Polyakov, X.-W. Zhu, L. Tian, P. Mukherjee, and X. Liu, "All-assay-max2 pqsar: Activity predictions as accurate as four-concentration ic50s for 8558 novartis assays," *Journal of Chemical Information and Modeling*, vol. 59, no. 10, pp. 4450–4459, 2019. PMID: 31518124.

[22] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," *arXiv preprint arXiv:1606.01540*, 2016.

[23] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *CoRR*, vol. abs/1707.06347, 2017.

[24] M. Ahn, H. Zhu, K. Hartikainen, H. Ponte, A. Gupta, S. Levine, and V. Kumar, "ROBEL: RObotics BEnchmarks for Learning with low-cost robots," in *Conference on Robot Learning (CoRL)*, 2019.

[25] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033, IEEE, 2012.

[26] A. Hill, A. Raffin, M. Ernestus, A. Gleave, A. Kanervisto, R. Traore, P. Dhariwal, C. Hesse, O. Klimov, A. Nichol, M. Plappert, A. Radford, J. Schulman, S. Sidor, and Y. Wu, "Stable baselines." `https://github.com/hill-a/stable-baselines`, 2018.

[27] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *ICML*, 2018.

Appendices

A Data Collection

In this section, we detail the data collection steps used for creating each of the tasks in design-bench. We answer (1) where is the data from, and (2) what pre-processing steps are used?

A.1 GFP-v0

The GFP task provided is a derivative of the GFP dataset [19]. The dataset we use in practice is that provided by [4] at the url https://github.com/dhbrookes/CbAS/tree/master/data. We process the dataset such that a single training example consists of a protein represented as a tensor $x_{\text{GFP}} \in \{0, 1\}^{238 \times 20}$. This tensor is a sequence of 238 one-hot vectors corresponding to which amino acid is present in that location in the protein. We use the dataset format of [4] with no additional processing. The data was originally collected by performing laboratory experiments constructing proteins similar to the Aequorea victoria green fluorescent protein and measuring fluorescence.

A.2 MoleculeActivity-v0

The MoleculeActivity task is a derivative of a much larger dataset that is derived from ChEMBL [16], a large database of chemicals and their properties. The data, similar to GFP, was originally collected by performing physical experiments on a large number of molecules, and measuring their activity with a target assay. We have processed the original dataset presented in [21], which consists of more than one million molecules and 11,000 assays, into a smaller scale task with 4216 molecules and a single assay. We select this assay by first calculating the number of unique scores present in the dataset per assay, and sorting the assays by the number of unique scores. We select assay 600885 from [21]. This particular assay has 4216 molecules after pre-processing. Our pre-processing steps include converting each molecule into a one-hot tensor $x_{\text{Molecule}} \in \{0, 1\}^{1024 \times 2}$. This is performed by calculating the Morgan radius 2 substructure fingerprints of 1024 bits, which is implemented in RDKit. This calculation requires the SMILES representation for each molecule, which is provided by [21]. The final step of pre-processing, is to sub sample the dataset by defining a percentile used to select and discard high-performing molecules, such that difficulty of the task is artificially increased. We use a split percentile of 80 for MoleculeActivity in the experiments in this paper.

A.3 Superconductor-v0

Superconductor-v0 is inspired by recent work [6] that applies offline MBO to optimize the properties of superconducting materials for high critical temperature. The data we provide in our benchmark is real-world superconductivity data originally collected by [2], and subsequently made available to the public at the url https://archive.ics.uci.edu/ml/datasets/Superconductivty+Data#. The original dataset consists of superconductors featurized into vectors $x_{\text{Superconductor}} \in \mathcal{R}^{81}$. One issue with the original dataset is that the largest value of a single dimension in the dataset is 22590.0, which appears to cause learning instability. We follow [6] and normalize each dimension of the design-space to have zero mean and unit variance. However, we deviate from the remaining pre-processing steps in [6]. In order to promote task realism, we directly use the superconductivity data, whereas [6] re-samples by collecting iid unit gaussian samples and labelling them with the task oracle function. This causes the scores in the dataset to correspond exactly to the scores provided by the oracle. No other domain in design-bench re-samples nor re-labels static data, so we omit it here for consistency.

A.4 HopperController-v0

The HopperController task is one that we provide ourselves. The goal of this task is to design a set of weights for as neural network policy, in order to achieve high expected return when evaluating that policy. The data collected for HopperController was taken by training a three layer neural network policy with 64 hidden units and 5126 total weights on the Hopper-v2 MuJoCo task using Proximal Policy Optimization [23]. Specifically, we use the default parameters for PPO provided in stable baselines [26]. The dataset we provide with this benchmark has 3200 unique weights. In order to collect this many, we run 32 experimental trials of PPO, where we train for one million steps, and save the weights of the policy every 10,000 environment steps. The policy weights are represented originally as a list of tensors. We first traverse this list and flatten each of the tensors, and we then concatenate each of these flattened tensors into a single training example $x_{\text{Hopper}} \in \mathcal{R}^{5126}$. The result is an optimization problem over neural network weights. After collecting these weights, we perform no additional pre-processing steps. In order to collect scores we perform a single rollout for each x using the Hopper-v2 MuJoCo environment. The horizon length for training and evaluation is limited to 1000 simulation time steps.

A.5 AntMorphology-v0 & DKittyMorphology-v0

Both morphology tasks are collected by us, and share methodology. The goal of these tasks is to design the morphology of a quadrupedal robot—an ant or a D'Kitty—such that the agent is able to crawl quickly in a particular direction. In order to collect data for this environment, we create variants of the MuJoCo Ant and the ROBEL D'Kitty agents that have parametric morphologies. The goal is to determine a mapping from the morphology of the agent to the average return that an agent trained for a particular intended morphology achieves. We implement this by pre-training a neural network policy using SAC [27]. For both the Ant and the D'Kitty, we train agents for up to three million environments steps, and a maximum episode length of 1000, with all other settings as default. These agents are pre-trained for a fixed gold-standard morphology—the default morphology of the Ant and D'Kitty respectively. Each morphology task consists of samples obtained by adding Gaussian noise with standard deviation 0.02 for Ant and 0.01 for DKitty times the design-space range to the gold-standard morphology. We label each sampled morphology by averaging the return of 16 rollouts of length 100 of an agent with that morphology.

B Oracle Functions

We detail oracle functions for evaluating ground truth scores for each of the tasks in design-bench. A common thread among these is that the oracle, if trained, is fit to a larger static dataset containing higher performing designs than observed by a downstream MBO algorithm.

B.1 GFP-v0

GFP-v0 uses the oracle function from [4]. This oracle is a Gaussian Process regression model with a protein-specific kernel proposed by [1]. The Gaussian Process is fit to a larger dataset than the static dataset packaged with GFP-v0, making it possible to sample a protein design that achieve a higher score than any other protein seen during training. The oracle score for GFP-v0 is implemented as the mean prediction of this Gaussian Process.

B.2 MoleculeActivity-v0

Following the procedure set by [21], the oracle function we use for MoleculeActivity-v0 is a random forest regression model. In particular, we use the RandomForestRegressor provided in scikit-learn, using identical hyperparameters to the random forest regression model used in [21]. The random forest is trained on the entire task dataset. In practice, samples that score at most the 80th percentile are observed by an MBO algorithm, which allows for sampling *unobserved* points that score higher than the highest training point.

B.3 Superconductor-v0

The Superconductor-v0 oracle function is also a random forest regression model. The model we use it the model described by [2]. We borrow the hyperparameters described by them, and we use the RandomForestRegressor provided in scikit-learn. Similar to the setup for the previous two tasks, this oracle is trained on the entire static dataset, and the task is instantiated with a split percentile. Samples scoring at most in the 80th percentile are observed by an MBO algorithm, which allows for sampling *unobserved* points that score in the unobserved top 20 percent.

B.4 HopperController-v0

Unlike the previously described tasks, HopperController-v0 and the remaining tasks implement an exact oracle function. For HopperController-v0 the oracle takes the form of a single rollout using the Hopper-v2 MuJoCo environment. The designs for HopperController-v0 are neural network weights, and during evaluation, a policy with those weights is instantiated—in this case that policy is a three layer neural network with 11 input units, two layers with 64 hidden units, and a final layer with 3 output units. The intermediate activations between layers are hyperbolic tangents. After building a policy, the Hopper-v2 environment is reset and the reward for 1000 time-steps is summed. That summed reward constitutes the score returned by the HopperController-v0 oracle. The limit of performance is the maximum return that an agent can achieve in Hopper-v2 over 1000 steps.

B.5 AntMorphology-v0 & DKittyMorphology-v0

The final two tasks in design-bench use an exact oracle function, using the MuJoCo simulator. For both morphology tasks, the simulator performs 16 rollouts and averages the sum of rewards attained over them. Each task is accompanied by a pre-trained neural network policy. To perform evaluation, a morphology is passed to the Ant or D'Kitty MuJoCo environments respectively, and a dynamic-morphology agent is initialized inside these environments. These environments are very sensitive to small morphological changes, and exhibit a high degree of stochasticity as a result. To compensate for the increased stochasticity, we average returns over 16 rollouts.