

ExPT: Scaling Foundation Models for Experimental Design via Synthetic Pretraining

Tung Nguyen
Sudhanshu Agrawal
Aditya Grover

University of California, Los Angeles

TUNGND@CS.UCLA.EDU
SUDHANSHUAGR27@G.UCLA.EDU
ADITYAG@CS.UCLA.EDU

Abstract

Experimental design is a fundamental problem in many science and engineering fields. In this problem, sample efficiency is crucial due to the time, money, and safety costs of real-world design evaluations. Existing approaches either rely on active data collection or access to large, labeled datasets of past experiments, making them impractical in many real-world scenarios. In this work, we address the more challenging yet realistic setting of *few-shot* experimental design, where only a few labeled data points of input designs and their corresponding values are available. We approach this problem as a conditional generation task, where a model conditions on a few labeled examples and the desired output to generate an optimal input design. To this end, we present **Pretrained Transformers for Experimental Design (ExPT)**, that uses a novel combination of synthetic pretraining with in-context learning to enable few-shot generalization. In ExPT, we only assume knowledge of a finite collection of unlabelled data points from the input domain and pretrain a transformer neural network to optimize diverse synthetic functions defined over this domain. Unsupervised pretraining allows ExPT to adapt to any design task at test time in an in-context fashion by conditioning on a few labeled data points from the target task and generating the candidate optima. We evaluate ExPT on few-shot experimental design in challenging domains and demonstrate its superior generality and performance compared to existing methods.

1. Introduction

Experimental design (ED) is ubiquitous in many fields of science and engineering, including material design (Hamidieh, 2018), protein (Brookes et al., 2019; Sarkisyan et al., 2016; Angermueller et al., 2020) or molecular (Gaulton et al., 2012) design, mechanical design (Berkenkamp et al., 2016; Liao et al., 2019), and neural architecture optimization (Zoph and Le, 2016). Sample efficiency is the most important criterion in these problems, as evaluating a design often involves expensive real-world experiments. The most common approach learns a surrogate to approximate the objective function and relies on active data collection via conducting online experiments to improve its estimate (Snoek et al., 2012). However, online data acquisition may be infeasible in the real world due to high costs, time constraints, or safety concerns. As an alternate, recent works have proposed offline ED (Brookes et al., 2019; Trabucco et al., 2022; Kumar and Levine, 2020; Trabucco et al., 2021; Chen et al., 2022; Krishnamoorthy et al., 2022), wherein a model learns to perform optimization from a fixed dataset of past experiments. While more practical than the online setting, current

offline methods and benchmarks assume access to large datasets of thousands of data points, which are hard or even impossible to obtain in high-stake and emerging science problems.

In this paper, we aim to overcome these limitations for hyper-efficient ED with only a handful of labeled examples. We introduce *few-shot* experimental design, a more challenging setting that better resembles real-world scenarios. In few-shot ED, we are not allowed to conduct additional experiments, and only have access to a few hundred labeled data points, as opposed to tens of thousands of data points in previous works. We approach few-shot ED with a pretraining-adaptation pipeline. In the pretraining phase, we only assume access to *unlabeled* data, i.e., input designs without associated function values. During adaptation, we have access to a few labeled examples of past experiments to adapt the model to the downstream task. This setup offers several advantages. First, it alleviates the requirement for costly annotated data and relies mainly on unlabeled inputs that are easily accessible. Second, unsupervised pretraining enables us to utilize the same pretrained backbone for adapting to multiple downstream optimization tasks within the same domain. For example, in molecule design, one may want to optimize for multiple properties, including drug-likeness, synthesizability, or similarity to target molecules (Brown et al., 2019; Gao et al., 2022).

The key question is how to make use of the unlabeled data to facilitate efficient generalization to downstream tasks during optimization. Our intuition is that, while the objective function is unknown, we can use the unlabeled inputs to generate pretraining data from *other* functions. If a model can few-shot learn from a diverse and challenging set of functions, it should be able to generalize quickly to any objective. This insight gives rise to our idea of *synthetic pretraining*, wherein we pretrain the model on data generated from a rich family of synthetic functions that operate on the same domain as the target task. Specifically, for each function drawn from this family, we sample a set of points by using the unlabeled data as inputs. We divide these points into a small context set and a target set, and train the model to perform conditional generation of the target input x given the context points and the target value y . A model that works well on this task should be able to efficiently capture the structures of the underlying function, i.e., how different regions of the input space influence the function value, from a small context set. By explicitly training the model to perform this task on a diverse set of functions, the model can generalize efficiently to downstream functions given limited supervision. After pretraining, we can perform optimization by simply conditioning the model on a few labeled examples from the downstream task and asking the model to generate an input that achieves the optimal score y^* .

We instantiate our model with a transformer-based architecture (Vaswani et al., 2017), which we call **P**retrained **T**ransformers for **E**xperimental Design (ExPT). ExPT is an encoder-decoder architecture, in which the encoder is a transformer (Vaswani et al., 2017) network that encodes the context points and the target value, and the decoder is a VAE (Kingma and Welling, 2013) model that predicts the high-dimensional target input. The transformer encoder allows ExPT to perform few-shot generation and optimization purely through in-context learning in a gradient-free fashion. We compare the performance of ExPT and various baselines on 2 few-shot settings created from Design-Bench (Trabucco et al., 2022), a standard database benchmark for ED. In both these settings, ExPT achieves the highest average score and the highest average ranking with respect to median performance, mean performance, and best-achieved performance. Especially in the more challenging setting, ExPT outperforms the second-best method by 70% in terms of the mean performance.

2. Synthetic Pretraining for Few-shot Experimental Design

2.1 Problem setup

Let $f : \mathcal{X} \rightarrow \mathbb{R}$ be a black-box function that operates on a d -dimensional domain $\mathcal{X} \subseteq \mathbb{R}^d$. In experimental design (ED), the goal is to find the input x^* that maximizes f : $x^* \in \arg \max_{x \in \mathcal{X}} f(x)$. We propose *few-shot ED*, a more challenging yet realistic setting to overcome these limitations. In few-shot ED, the goal is to optimize *any* objective function in the domain \mathcal{X} given only a handful of examples. We approach this problem with a pretraining-adaptation pipeline. In the *pretraining* phase, we assume access to an *unlabeled* dataset $\mathcal{D}_{\text{unlabeled}} = \{x_i\}_{i=1}^{|\mathcal{D}_{\text{unlabeled}}|}$ from the optimization domain $\mathcal{X} \subseteq \mathbb{R}^d$.

During the *adaptation* phase, one can use the pretrained model to optimize any objective function f in the same domain \mathcal{X} . We now have access to a few-shot *labeled* dataset to adapt to the downstream function $\mathcal{D}_{\text{few-shot}} = \{(x_1, y_1), \dots, (x_{|\mathcal{D}_{\text{few-shot}}|}, y_{|\mathcal{D}_{\text{few-shot}}|})\}$, in which $y_i = f(x_i)$ and $|\mathcal{D}_{\text{few-shot}}| \ll |\mathcal{D}_{\text{unlabeled}}|$. After adaptation, we evaluate a few-shot optimization method by allowing it to propose Q input x 's and query their scores using the black-box function f , where Q is often called the optimization budget (Trabucco et al., 2022; Kumar and Levine, 2020; Trabucco et al., 2021; Chen et al., 2022; Krishnamoorthy et al., 2022). The performance of a black-box optimizer is then measured by computing the max, median, and mean of the Q evaluations, This setup provides 2 key benefits. First, it resembles many real-world scenarios, where the potential design inputs are cheap and easy to obtain while their target function values are expensive to evaluate. For example, in molecular optimization, we have databases of millions of molecules (Kim et al., 2016; Blum and Raymond, 2009; Ruddigkeit et al., 2012) but only the properties of a handful are known (Irwin and Shoichet, 2005; Ramakrishnan et al., 2014; Gaulton et al., 2017). Second, unsupervised pretraining allows us to train a general backbone that we can adapt to multiple optimization tasks in the same domain.

2.2 Synthetic Pretraining for Few-shot ED

Intuitively, the adaptation phase in §2.1 resembles a few-shot learning problem, in which a model is tasked to produce the optimal input x^* by conditioning on a few labeled examples in $\mathcal{D}_{\text{few-shot}}$. To perform well in this task, a model has to efficiently capture the structure of a high-dimension function f , i.e., what regions of the function lead to higher values and vice versa, from very few examples in $\mathcal{D}_{\text{few-shot}}$. Given this perspective, the question now is how to make use of the unlabeled dataset $\mathcal{D}_{\text{unlabeled}}$ to pretrain a model that achieves such efficient generalization to the objective function f . Our key insight is, if a model learns to perform few-shot learning on a diverse and challenging set of functions, it should be able to adapt quickly to any objective function at test time. While the function values are unknown during pretraining, we can use the unlabeled inputs x 's to generate pretraining data from *other* functions. This gives rise to our idea of *synthetic pretraining*, wherein we pretrain the model to perform few-shot learning on a family of synthetic functions \tilde{F} that operate on the same input domain \mathcal{X} of the objective f . We discuss in detail our mechanism for synthetic data generation in Section 3.1. For each function \tilde{f} generated from \tilde{F} , we sample a set of function evaluations $\{(x_i, y_i)\}_{i=1}^N$ divided into a small context set $\{(x_i, y_i)\}_{i=1}^m$ and a target set $\{(x_j, y_j)\}_{j=m+1}^N$. The model predicts the target points conditioning on the context set.

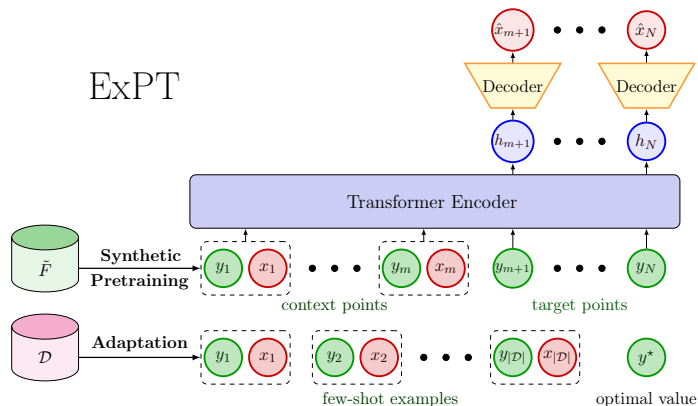


Figure 1: The pretraining-adaptation approach and ExPT architecture. We sample synthetic data from \tilde{F} and pretrain the model to maximize $\log p(x_{m+1:N} | x_{1:m}, y_{1:m}, y_{m+1:N})$. At adaptation, the model conditions on $\mathcal{D}_{\text{few-shot}}$ and y^* to generate candidates.

There are two different approaches to pretraining a model on this synthetic data. The first approach is forward modeling, where the model predicts the target outputs $y_{m+1:N}$ given the context points and the target inputs $x_{m+1:N}$. During adaptation, we can condition the model on the labeled examples in $\mathcal{D}_{\text{few-shot}}$ and perform gradient ascent updates to improve an existing design input x_t . However, as commonly observed in previous works (Trabucco et al., 2022; Kumar and Levine, 2020; Trabucco et al., 2021), this approach is susceptible to producing highly suboptimal inputs. This is because performing gradient ascent with respect to an imperfect forward model may result in points that have high values under the model but are poor when evaluated using the real function. Instead, we propose to perform *inverse modeling*, where the model learns to predict the inputs $x_{m+1:N}$ given the output values $y_{m+1:N}$ and the context points. As the model learns to directly generate input x 's, it is less vulnerable to the aforementioned problem. Figure 1 illustrates the proposed pretraining and adaptation pipeline.

After pretraining, ExPT can adapt to any objective f in the domain in a gradient-free fashion. Samples in the few-shot dataset $\mathcal{D}_{\text{few-shot}}$ become the context points and the model conditions on only one target y^* , which is the optimal value of f , to generate candidate optima. Note that we only assume the knowledge of y^* and not x^* . This assumption is common in many prior works (Krishnamoorthy et al., 2022; Nguyen and Osborne, 2020; Chen et al., 2021; Emmons et al., 2021). In practice, y^* might be known based on domain knowledge. For example, in molecule design, there are physical limits on the value of certain properties such as relaxed energy, in robot applications, the optimal performance can be computed from the cost function, and in neural architecture search, we know the theoretical limits on the highest possible accuracy for classifiers.

3. Pretrained Transformers for Experimental Design

Section 2 presents our general idea of synthetic pretraining for few-shot ED. In this section, we present the details of synthetic data generation and our proposed model architecture, the two components that constitute **P**retrained **T**ransformers for **E**xperimental **D**esign (ExPT).

3.1 Data generation

We need a family of functions to generate synthetic data for pretraining ExPT. A good family of functions should be easy to sample from and should be capable of producing diverse functions. We choose to generate synthetic data from Gaussian Processes with an RBF kernel. This is for several reasons. First, they are a natural choice as they represent distributions over functions. Second, it is easy and cheap to sample data from prior GPs. And third, a GP with an RBF kernel is a universal approximator to any function (Micchelli et al., 2006). Specifically, $\tilde{f} \sim \mathcal{GP}(0, \mathcal{K})$, $\mathcal{K}(x, x') = \sigma^2 \exp\left(-\frac{(x-x')^2}{2\ell^2}\right)$, in which σ and ℓ are the two hyperparameters of the RBF kernel. In practice, we randomize both σ and ℓ to increase the diversity of the pretraining data. Appendix E.1 demonstrates the empirical importance of these hyperparameters.

3.2 Model architecture

We need a model architecture that can condition on a few examples drawn from an underlying function to make predictions for other points. This resembles the idea of *in-context learning* that has proven very successful in language (Brown et al., 2020; Min et al., 2022) and other domains (Müller et al., 2021; Nguyen and Grover, 2022; Garg et al., 2022; Laskin et al.). The key to success in these works is a transformer architecture that performs in-context learning efficiently via the attention mechanism (Vaswani et al., 2017). Inspired by this, we instantiate ExPT with a transformer-based architecture. Figure 1 illustrates the ExPT overall architecture. Specifically, ExPT employs a transformer encoder that encodes the context points $\{(x_i, y_i)\}_{i=1}^m$ and the target inputs $y_{m+1:N}$, and outputs hidden vectors $h_{m+1:N}$. To inform the model that x_i and y_i are the input and the corresponding output from \tilde{f} , we concatenate them to form a token. This results in the sequence $\{(y_1, x_1), \dots, (y_m, x_m), y_{m+1}, \dots, y_N\}$. We then embed these tokens using two 1-layer MLP networks, one for the pairs and one for the target inputs, before feeding the sequence to the transformer layers. We implement a masking mechanism that prevents the context points from attending the target points, as they do not contain information about the underlying function \tilde{f} .

Each hidden vector h_i output by the transformer encoder encompasses the information of the context points and the target input y_i . Therefore, given h_i , the conditional probability $p_\theta(x_i | x_{1:m}, y_{1:m}, y_i)$ reduces to $p_\theta(x_i | h_i)$. As x_i is high-dimensional, we can utilize existing generative models to model the conditional distribution $p(x_i | h_i)$. In this work, we train a conditional VAE model (Kingma and Welling, 2013) alongside the transformer encoder because of its training stability, light hyperparameter tuning, and good empirical performance. For discrete tasks, we follow the same procedure as Trabucco et al. (2022) that emulates logit values by interpolating between a uniform distribution and the one hot values. We train the entire model by maximizing the lower bound of the conditional likelihood $\log p(x_i | h_i)$.

4. Experiments

4.1 Design-Bench experiments

Tasks We consider 4 tasks from Design-Bench (Trabucco et al., 2022). In D’kitty and Ant, the goal is to optimize the morphological structure of two simulated robots, Ant (Brockman

Table 1: Comparison of ExPT and the baselines on the few-shot **random** setting of 5 Design-Bench tasks. We report median, max, and mean performance across 3 random seeds. Higher scores and lower ranks are better. **Blue** denotes the best entry in the column, and **Violet** denotes the second best.

	Baseline	D’Kitty	Ant	TF Bind 8	TF Bind 10	Mean score (\uparrow)	Mean rank (\downarrow)
	$\mathcal{D}_{\text{few-shot}}(\text{best})$	0.883	0.563	0.439	0.466	—	—
Median	MINs	0.859 \pm 0.014	0.485 \pm 0.152	0.416 \pm 0.019	0.468 \pm 0.014	0.557 \pm 0.050	4.0
	COMs	0.752 \pm 0.007	0.411 \pm 0.012	0.371 \pm 0.001	0.468 \pm 0.000	0.501 \pm 0.005	4.0
	BONET	0.852 \pm 0.013	0.597 \pm 0.119	0.441 \pm 0.003	0.483 \pm 0.009	0.593 \pm 0.036	2.3
	BDI	0.592 \pm 0.020	0.396 \pm 0.018	0.540 \pm 0.032	0.438 \pm 0.034	0.492 \pm 0.026	4.8
	GP-qEI	0.842 \pm 0.058	0.550 \pm 0.007	0.439 \pm 0.000	0.467 \pm 0.000	0.575 \pm 0.016	4.0
	ExPT	0.902 \pm 0.006	0.705 \pm 0.018	0.473 \pm 0.014	0.477 \pm 0.014	0.639 \pm 0.013	1.5
Max	MINs	0.930 \pm 0.010	0.890 \pm 0.017	0.814 \pm 0.030	0.639 \pm 0.017	0.818 \pm 0.019	3.3
	COMs	0.920 \pm 0.010	0.841 \pm 0.044	0.686 \pm 0.152	0.656 \pm 0.020	0.776 \pm 0.057	4.0
	BONET	0.909 \pm 0.012	0.888 \pm 0.024	0.887 \pm 0.053	0.702 \pm 0.006	0.847 \pm 0.024	3.0
	BDI	0.918 \pm 0.006	0.806 \pm 0.094	0.906 \pm 0.074	0.532 \pm 0.023	0.791 \pm 0.049	4.5
	GP-qEI	0.896 \pm 0.000	0.887 \pm 0.000	0.513 \pm 0.104	0.647 \pm 0.011	0.736 \pm 0.029	5.0
	ExPT	0.973 \pm 0.005	0.970 \pm 0.004	0.933 \pm 0.036	0.677 \pm 0.048	0.888 \pm 0.023	1.3
Mean	MINs	0.624 \pm 0.025	0.009 \pm 0.013	0.415 \pm 0.030	0.465 \pm 0.015	0.378 \pm 0.021	4.8
	COMs	0.515 \pm 0.050	0.020 \pm 0.006	0.369 \pm 0.003	0.471 \pm 0.004	0.344 \pm 0.016	4.8
	BONET	0.837 \pm 0.023	0.579 \pm 0.024	0.448 \pm 0.011	0.484 \pm 0.009	0.587 \pm 0.017	2.0
	BDI	0.570 \pm 0.032	0.385 \pm 0.012	0.536 \pm 0.032	0.444 \pm 0.027	0.484 \pm 0.026	3.5
	GP-qEI	0.505 \pm 0.006	0.019 \pm 0.001	0.439 \pm 0.001	0.473 \pm 0.002	0.359 \pm 0.003	4.5
	ExPT	0.865 \pm 0.016	0.639 \pm 0.026	0.476 \pm 0.010	0.474 \pm 0.015	0.614 \pm 0.017	1.5

et al., 2016) to run as fast as possible, and D’kitty (Ahn et al., 2020) to reach a fixed target location. **TF Bind 8** and **TF Bind 10** are two discrete tasks, where the goal is to find the length-8 and length-10 DNA sequence that has a maximum binding affinity with the SIX6_REF_R1 transcription factor. For each task, Design-Bench provides a **public** dataset, a larger **hidden** dataset which is used to normalize the scores, and an oracle. We have an exact oracle to evaluate the proposed designs in all 4 tasks we consider.

Few-shot settings We create 2 few-shot settings from the above tasks, which we call **random** and **poorest**. In **random**, we randomly subsample 1% of data points in the **public** set of each task as the few-shot dataset $\mathcal{D}_{\text{few-shot}}$. The **poorest** setting is more challenging, where we use 1% of the data points that have the *lowest* scores. Because of the limited space, we only present the **random** setting in the main paper.

Baselines We compare ExPT with BayesOpt (GP-qEI) (Snoek et al., 2012), a canonical ED method, and MINs Kumar and Levine (2020), COMs Trabucco et al. (2021), BDIChen et al. (2022), and BONET Krishnamoorthy et al. (2022), four recent deep learning models that have achieved state-of-the-art performance in the offline setting.

Evaluation For each considered method, we allow an optimization budget $Q = 256$. We report the median score, the max score, and the mean score among the 256 proposed inputs. Following previous works, we normalize the score to $[0, 1]$ by using the minimum and maximum function values from a large hidden dataset $y_{\text{norm}} = \frac{y - y_{\text{min}}}{y_{\text{max}} - y_{\text{min}}}$. We report the mean and standard deviation of the score across 3 independent runs for each method.

Results Table 1 shows the performance of different methods in the **random** setting. Most methods perform well in the **random** setting, where ExPT achieves the highest average score and the best average rank across all 3 performance metrics. For each of the tasks and metrics considered, ExPT is either the best or second-best performing method. Notably, in Ant, ExPT significantly outperforms the best baseline by 18%, 9%, and 10% with respect to the median, max, and mean performance, respectively. Only ExPT and BONET achieve a meaningful performance in Ant when considering the mean score. BONET is also the overall second-best method in this setting.

References

- Michael Ahn, Henry Zhu, Kristian Hartikainen, Hugo Ponte, Abhishek Gupta, Sergey Levine, and Vikash Kumar. Robel: Robotics benchmarks for learning with low-cost robots. In *Conference on robot learning*, pages 1300–1313. PMLR, 2020.
- Christof Angermueller, David Dohan, David Belanger, Ramya Deshpande, Kevin Murphy, and Lucy Colwell. Model-based reinforcement learning for biological sequence design. 2020.
- Felix Berkenkamp, Angela P Schoellig, and Andreas Krause. Safe controller optimization for quadrotors with gaussian processes. In *2016 IEEE international conference on robotics and automation (ICRA)*, pages 491–496. IEEE, 2016.
- Lorenz C Blum and Jean-Louis Reymond. 970 million druglike small molecules for virtual screening in the chemical universe database gdb-13. *Journal of the American Chemical Society*, 131(25):8732–8733, 2009.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- David Brookes, Hahnbeom Park, and Jennifer Listgarten. Conditioning by adaptive sampling for robust design. In *International conference on machine learning*, pages 773–782. PMLR, 2019.
- Nathan Brown, Marco Fiscato, Marwin HS Segler, and Alain C Vaucher. Guacamol: benchmarking models for de novo molecular design. *Journal of chemical information and modeling*, 59(3):1096–1108, 2019.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020.
- Kaidi Cao, Jiakuan You, Jiaju Liu, and Jure Leskovec. Autotransfer: AutoML with knowledge transfer - an application to graph neural networks. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=y81ppNf_vg.
- Can Chen, Yingxueff Zhang, Jie Fu, Xue Steve Liu, and Mark Coates. Bidirectional learning for offline infinite-width model-based optimization. *Advances in Neural Information Processing Systems*, 35:29454–29467, 2022.
- Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097, 2021.

- Scott Emmons, Benjamin Eysenbach, Ilya Kostrikov, and Sergey Levine. Rvs: What is essential for offline rl via supervised learning? *arXiv preprint arXiv:2112.10751*, 2021.
- Clara Fannjiang and Jennifer Listgarten. Autofocused oracles for model-based design. *Advances in Neural Information Processing Systems*, 33:12945–12956, 2020.
- Justin Fu and Sergey Levine. Offline model-based optimization via normalized maximum likelihood estimation. *arXiv preprint arXiv:2102.07970*, 2021.
- Wenhao Gao, Tianfan Fu, Jimeng Sun, and Connor Coley. Sample efficiency matters: a benchmark for practical molecular optimization. *Advances in Neural Information Processing Systems*, 35:21342–21357, 2022.
- Victor Garcia and Joan Bruna. Few-shot learning with graph neural networks, 2018.
- Shivam Garg, Dimitris Tsipras, Percy S Liang, and Gregory Valiant. What can transformers learn in-context? a case study of simple function classes. *Advances in Neural Information Processing Systems*, 35:30583–30598, 2022.
- Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo Rezende, and SM Ali Eslami. Conditional neural processes. In *International conference on machine learning*, pages 1704–1713. PMLR, 2018a.
- Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J Rezende, SM Eslami, and Yee Whye Teh. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018b.
- Anna Gaulton, Louisa J Bellis, A Patricia Bento, Jon Chambers, Mark Davies, Anne Hersey, Yvonne Light, Shaun McGlinchey, David Michalovich, Bissan Al-Lazikani, et al. ChEMBL: a large-scale bioactivity database for drug discovery. *Nucleic acids research*, 40(D1): D1100–D1107, 2012.
- Anna Gaulton, Anne Hersey, Michał Nowotka, A Patricia Bento, Jon Chambers, David Mendez, Prudence Mutowo, Francis Atkinson, Louisa J Bellis, Elena Cibrián-Uhalte, et al. The ChEMBL database in 2017. *Nucleic acids research*, 45(D1):D945–D954, 2017.
- Ethan Goan and Clinton Fookes. Bayesian neural networks: An introduction and survey. *Case Studies in Applied Bayesian Data Science: CIRM Jean-Morlet Chair, Fall 2018*, pages 45–87, 2020.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- Jonathan Gordon, Wessel P Bruinsma, Andrew YK Foong, James Requeima, Yann Dubois, and Richard E Turner. Convolutional conditional neural processes. In *International Conference on Learning Representations*.
- Kam Hamidieh. A data-driven statistical model for predicting the critical temperature of a superconductor. *Computational Materials Science*, 154:346–354, 2018.

- John J Irwin and Brian K Shoichet. Zinc- a free database of commercially available compounds for virtual screening. *Journal of chemical information and modeling*, 45(1):177–182, 2005.
- Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, Ali Eslami, Dan Rosenbaum, Oriol Vinyals, and Yee Whye Teh. Attentive neural processes. In *International Conference on Learning Representations*.
- Sunghwan Kim, Paul A Thiessen, Evan E Bolton, Jie Chen, Gang Fu, Asta Gindulyte, Lianyi Han, Jane He, Siqian He, Benjamin A Shoemaker, et al. Pubchem substance and compound databases. *Nucleic acids research*, 44(D1):D1202–D1213, 2016.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Kundan Krishna, Jeffrey Bigham, and Zachary C Lipton. Does pretraining for summarization require knowledge transfer? *arXiv preprint arXiv:2109.04953*, 2021.
- Siddarth Krishnamoorthy, Satvik Mehul Mashkaria, and Aditya Grover. Generative pretraining for black-box optimization. *arXiv preprint arXiv:2206.10786*, 2022.
- Aviral Kumar and Sergey Levine. Model inversion networks for model-based optimization. *Advances in Neural Information Processing Systems*, 33:5126–5137, 2020.
- Michael Laskin, Luyu Wang, Junhyuk Oh, Emilio Parisotto, Stephen Spencer, Richie Steigerwald, DJ Strouse, Steven Stenberg Hansen, Angelos Filos, Ethan Brooks, et al. In-context reinforcement learning with algorithm distillation. In *NeurIPS 2022 Foundation Models for Decision Making Workshop*.
- Thomas Liao, Grant Wang, Brian Yang, Rene Lee, Kristofer Pister, Sergey Levine, and Roberto Calandra. Data-efficient learning of morphology and controller for a microrobot. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2488–2494. IEEE, 2019.
- Daniel James Lizotte. Practical bayesian optimization. 2008.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- Charles A Micchelli, Yuesheng Xu, and Haizhang Zhang. Universal kernels. *Journal of Machine Learning Research*, 7(12), 2006.
- Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hannaneh Hajishirzi. Metaicl: Learning to learn in context. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2791–2809, 2022.

- Samarth Mishra, Rameswar Panda, Cheng Perng Phoo, Chun-Fu (Richard) Chen, Leonid Karlinsky, Kate Saenko, Venkatesh Saligrama, and Rogerio S. Feris. Task2sim: Towards effective pre-training and transfer from synthetic data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9194–9204, June 2022.
- Samuel Müller, Noah Hollmann, Sebastian Pineda Arango, Josif Grabocka, and Frank Hutter. Transformers can do bayesian inference. *arXiv preprint arXiv:2112.10510*, 2021.
- Tung Nguyen and Aditya Grover. Transformer neural processes: Uncertainty-aware meta learning via sequence modeling. *arXiv preprint arXiv:2207.04179*, 2022.
- Vu Nguyen and Michael A. Osborne. Knowing the what but not the where in bayesian optimization, 2020.
- Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole Von Lilienfeld. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1): 1–7, 2014.
- Lars Ruddigkeit, Ruud Van Deursen, Lorenz C Blum, and Jean-Louis Reymond. Enumeration of 166 billion organic small molecules in the chemical universe database gdb-17. *Journal of chemical information and modeling*, 52(11):2864–2875, 2012.
- Karen S Sarkisyan, Dmitry A Bolotin, Margarita V Meer, Dinara R Usmanova, Alexander S Mishin, George V Sharonov, Dmitry N Ivankov, Nina G Bozhanova, Mikhail S Baranov, Onuralp Soylemez, et al. Local fitness landscape of the green fluorescent protein. *Nature*, 533(7603):397–401, 2016.
- Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015.
- Gautam Singh, Jaesik Yoon, Youngsung Son, and Sungjin Ahn. Sequential neural processes. *Advances in Neural Information Processing Systems*, 32, 2019.
- Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30, 2017.
- Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical bayesian optimization of machine learning algorithms, 2012.
- Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.
- Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip H.S. Torr, and Timothy M. Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.

- Kevin Swersky, Jasper Snoek, and Ryan P Adams. Multi-task bayesian optimization. *Advances in neural information processing systems*, 26, 2013.
- Brandon Trabucco, Aviral Kumar, Xinyang Geng, and Sergey Levine. Conservative objective models for effective offline model-based optimization. In *International Conference on Machine Learning*, pages 10358–10368. PMLR, 2021.
- Brandon Trabucco, Xinyang Geng, Aviral Kumar, and Sergey Levine. Design-bench: Benchmarks for data-driven offline model-based optimization. In *International Conference on Machine Learning*, pages 21658–21676. PMLR, 2022.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Yaqing Wang, Quanming Yao, James T Kwok, and Lionel M Ni. Generalizing from a few examples: A survey on few-shot learning. *ACM computing surveys (csur)*, 53(3):1–34, 2020.
- Martin Wistuba and Josif Grabocka. Few-shot bayesian optimization with deep kernel surrogates. *arXiv preprint arXiv:2101.07667*, 2021.
- Chenfei Wu, Jian Liang, Lei Ji, Fan Yang, Yuejian Fang, Daxin Jiang, and Nan Duan. Nüwa: Visual synthesis pre-training for neural visual world creation. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVI*, pages 720–736. Springer, 2022.
- Yuhuai Wu, Markus N Rabe, Wenda Li, Jimmy Ba, Roger B Grosse, and Christian Szegedy. Lime: Learning inductive bias for primitives of mathematical reasoning. In *International Conference on Machine Learning*, pages 11251–11262. PMLR, 2021.
- Sihyun Yu, Sungsoo Ahn, Le Song, and Jinwoo Shin. Roma: Robust model adaptation for offline model-based optimization. *Advances in Neural Information Processing Systems*, 34: 4619–4631, 2021.
- Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016.

Appendix A. Related work

Online ED The majority of existing approaches solve ED in an active setting, where the model is allowed to query the black-box function to collect more data. Many of these works are based on Bayesian Optimization Lizotte (2008); Nguyen and Osborne (2020); Shahriari et al. (2015); Snoek et al. (2012); Swersky et al. (2013), which typically employs a surrogate model to the black-box function and an acquisition function. The surrogate model is often a predictive model that can quantify uncertainty, such as Gaussian Processes (Srinivas et al., 2009), Neural Processes (Garnelo et al., 2018a,b; Gordon et al.; Kim et al.; Singh et al., 2019; Nguyen and Grover, 2022), or Bayesian Neural Networks (Goan and Fookes, 2020). The acquisition function uses the uncertainty output by the surrogate model to trade off between exploration and exploitation for querying new points.

Offline ED Recent works have proposed to solve ED by learning from a fixed set of (x, y) pairs to bypass active data collection (Trabucco et al., 2022; Kumar and Levine, 2020; Trabucco et al., 2021; Chen et al., 2022; Krishnamoorthy et al., 2022; Fu and Levine, 2021; Brookes et al., 2019; Fannjiang and Listgarten, 2020; Yu et al., 2021). The Design-Bench benchmark Trabucco et al. (2022) consists of several such tasks in the physical sciences and robotics and is used by many recent works in offline ED. MINs Kumar and Levine (2020) and BONET Krishnamoorthy et al. (2022) perform optimization by generating designs x via conditioning on a high score value y . MINs uses a GANs (Goodfellow et al., 2020) model on (x, y) pairs and BONET casts offline ED as a sequence modeling problem. COMs Trabucco et al. (2021) formulates a conservative objective function that penalizes high-scoring poor designs and uses it to train a surrogate forward model which is then optimized using gradient ascent. BDI Chen et al. (2022) uses a bidirectional model consisting of a forward and backward models that learn mappings from the dataset to high-scoring designs and vice versa. In contrast to these works, we propose ExPT in the few-shot ED setting, where the model is given access to only the x 's during pretraining, and a handful of labeled examples for adaptation.

Synthetic Pretraining In the absence of vast amounts of labeled data, pretraining on synthetic data is an effective method for achieving significant gains in model performance. Prior works in this direction construct synthetic tasks which improve performance on diverse downstream tasks such as mathematical reasoning Wu et al. (2021), text summarization Krishna et al. (2021), and perception tasks in vision Mishra et al. (2022). Each synthetic task produces a dataset of labeled (x, y) values that can be used to train a model as usual for various objectives. Often, pre-training in this manner produces better results than simply pre-training on another real dataset. In this work, we demonstrate that pretraining on synthetic data generated from GPs can achieve significant generalization to downstream functions, leading to state-of-the-art performance on challenging few-shot optimization problems.

Few-shot learning Few-shot learning is a common paradigm in deep learning, where the model is pretrained on large amounts of data in an unsupervised manner. At test time, the model is given only a few examples from a downstream task and is expected to generalize Wang et al. (2020). This technique has found applications in text-generation (GPT-x) Brown et al. (2020), image classification (Sung et al., 2018; Snell et al., 2017), graph neural networks Garcia and Bruna (2018), text to visual-data generation Wu et al. (2022), and neural architecture search Wistuba and Grabocka (2021) Cao et al. (2023). ExPT is

capable of performing few-shot learning for black-box optimization in a variety of domains. Moreover, ExPT is pretrained on synthetically generated data with no prior knowledge of the downstream objective.

Appendix B. Additional experiments

B.1 Evaluation in the few-shot poorest setting

Table 2: Comparison of ExPT and the baselines on the few-shot **poorest** setting of 5 Design-Bench tasks. We report the median, max, and mean performance across 3 random seeds. Higher scores and lower ranks are better. **Blue** denotes the best entry in the column, and **Violet** denotes the second best.

	Baseline	D’Kitty	Ant	TF Bind 8	TF Bind 10	Mean score (\uparrow)	Mean rank (\downarrow)
	$\mathcal{D}_{\text{few-shot}}(\text{best})$	0.307	0.124	0.124	0.000	—	—
Median	MINs	0.480 \pm 0.156	0.316 \pm 0.040	0.437 \pm 0.007	0.463 \pm 0.003	0.424 \pm 0.052	3.5
	COMs	0.733 \pm 0.023	0.401 \pm 0.026	0.111 \pm 0.000	0.459 \pm 0.006	0.426 \pm 0.014	4.3
	BONET	0.310 \pm 0.000	0.236 \pm 0.047	0.319 \pm 0.018	0.461 \pm 0.017*	0.332 \pm 0.021	4.8
	BDI	0.309 \pm 0.000	0.192 \pm 0.012	0.365 \pm 0.000	0.454 \pm 0.017	0.330 \pm 0.007	5.5
	GP-qEI	0.883 \pm 0.000	0.565 \pm 0.001	0.439 \pm 0.000	0.467 \pm 0.000	0.589 \pm 0.000	2.0
	ExPT	0.922 \pm 0.009	0.686 \pm 0.090	0.552 \pm 0.042	0.489 \pm 0.013	0.662 \pm 0.039	1.0
Max	MINs	0.841 \pm 0.014	0.721 \pm 0.031	0.962 \pm 0.019	0.648 \pm 0.025	0.793 \pm 0.022	3.3
	COMs	0.931 \pm 0.022	0.843 \pm 0.020	0.124 \pm 0.000	0.739 \pm 0.057	0.659 \pm 0.025	3.3
	BONET	0.929 \pm 0.031	0.557 \pm 0.118	0.809 \pm 0.038	0.519 \pm 0.039*	0.704 \pm 0.057	5.0
	BDI	0.939 \pm 0.002	0.693 \pm 0.109	0.913 \pm 0.000	0.596 \pm 0.020	0.785 \pm 0.033	3.5
	GP-qEI	0.896 \pm 0.000	0.887 \pm 0.000	0.439 \pm 0.000	0.645 \pm 0.021	0.717 \pm 0.005	3.8
	ExPT	0.946 \pm 0.018	0.965 \pm 0.004	0.873 \pm 0.035	0.615 \pm 0.022	0.850 \pm 0.020	2.3
Mean	MINs	0.623 \pm 0.051	0.015 \pm 0.017	0.464 \pm 0.009	0.463 \pm 0.002	0.391 \pm 0.020	3.3
	COMs	0.607 \pm 0.021	0.033 \pm 0.003	0.109 \pm 0.001	0.454 \pm 0.004	0.301 \pm 0.007	4.8
	BONET	0.490 \pm 0.023	0.234 \pm 0.052	0.318 \pm 0.018	0.459 \pm 0.018	0.375 \pm 0.028	4.0
	BDI	0.364 \pm 0.004	0.215 \pm 0.021	0.369 \pm 0.000	0.453 \pm 0.018	0.350 \pm 0.011	4.8
	GP-qEI	0.533 \pm 0.001	0.018 \pm 0.000	0.439 \pm 0.000	0.470 \pm 0.002	0.365 \pm 0.001	3.5
	ExPT	0.871 \pm 0.018	0.646 \pm 0.061	0.549 \pm 0.032	0.488 \pm 0.011	0.639 \pm 0.031	1.0

Table 2 shows the superior performance of ExPT in few-shot **poorest**, the more challenging setting. ExPT achieves the highest score in 10/12 individual tasks and metrics, and also achieves the highest score and the best rank across tasks on average. Notably, in terms of the mean score, ExPT beats the best baseline by a large margin, achieving an improvement of 40%, 176%, and 18% on D’Kitty, Ant, and TF Bind 8, and 70% on average. The performance of most baselines drops significantly from the **random** to the **poorest** setting, including BONET, the second-best method in the **random** setting. This was also previously observed in the BONET paper (Krishnamoorthy et al., 2022). Interestingly, the performance of ExPT, MINs, and GP-qEI is not affected much by the quality of the few-shot data, and even improves in certain metrics. We hypothesize that even though the dataset is of lower quality, it may contain specific anti-correlation patterns about the problem that the model can exploit.

B.2 Synthetic experiments

We evaluate the performance of ExPT in a synthetic experiment, where we train ExPT on data generated from Gaussian Processes (GPs) with an RBF kernel and test the model on four out-of-distribution functions drawn from four different kernels: Matern, Linear, Cosine,

and Periodic. We pretrain the model for 2000 iterations with 128 synthetic functions at each iteration. We randomize the length scale parameter $\ell \sim \mathcal{U}[5.0, 10.0]$ and function scale parameter $\sigma_y \sim \mathcal{U}[1.0, 10.0]$ of the RBF kernel to increase pretraining data diversity. For each function generated, we sample 228 data points that we separate into 100 context points and 128 target points and train the model using the loss function in (??). Each input x is a 32-dimensional vector, and each dimension is sampled from a uniform distribution $\mathcal{U}[-3, 3]$. For each test function, we sample a large dataset of 20000 data points. We then randomly select 100 samples from the data points with function values lower than the 20th percentile as the few-shot data. We condition the model on this few-shot dataset and the maximal value y^* in the large dataset to generate 256 candidates and report the best score achieved. We normalize the score to $[0, 1]$ using the worst and the best value in the large dataset.

Figure 2 shows the performance of ExPT on four test functions through the course of training. The model performs well on all four functions, achieving scores that are much higher than the max value in the few-shot dataset, and approaching the true optimal value, even for kernels that are significantly different from RBF like Cosine and Periodic. Moreover, the performance improves consistently as we pretrain, showing that pretraining facilitates generalization to out-of-distribution functions with very few examples. See Appendix B for a detailed setup of this experiment.




Figure 2: The performance of ExPT on 4 out-of-distribution synthetic tasks through the pretraining phase. We average the performance across 3 seeds.

B.3 Pretraining analysis

In addition to the absolute performance, we investigate the performance of ExPT on downstream tasks through the course of pretraining. Figure 3 shows that the performance of ExPT in most tasks improves consistently as the number of pretraining steps increases. This shows that synthetic pretraining on diverse functions facilitates the generalization to complex real-world functions. In Ant, the performance slightly drops between 4000 and 10000 iterations. Therefore, we can further improve ExPT if we have a way to stop pretraining at a point that likely leads to the best performance in downstream tasks. However, in practice, we do not have the luxury of testing on real functions during pretraining. Alternatively, we could perform validation and early stopping on a set of held-out, out-of-distribution synthetic functions. We leave this to future work.

B.4 Few-shot optimization for multiple objectives

As we mention in Section 2, one advantage of unsupervised pretraining is the ability to optimize for multiple objectives during the adaptation phase. In this section, we show that

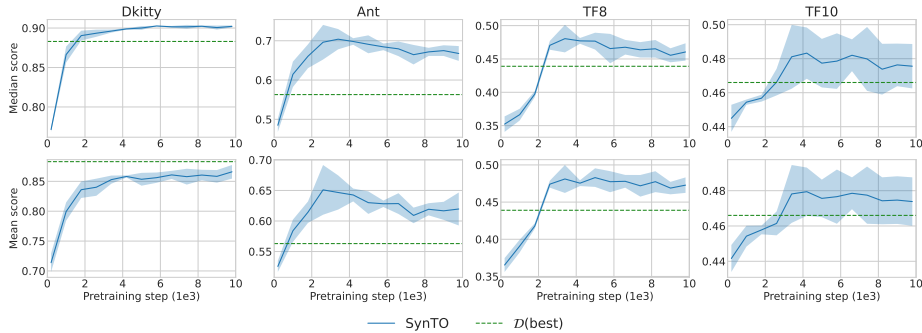


Figure 3: The median and mean performance of ExPT of 4 Design-Bench tasks through the course of pretraining. We average the performance across 3 seeds.

the same pretrained ExPT model is capable of optimizing different objectives in D’Kitty and Ant domains. We create two variants of the original D’Kitty, namely D’Kitty-45 and D’Kitty-60, whose objectives are to navigate the robot to goals that are 45° and 60° away from the original goal, respectively. For Ant, we create Ant- v_y where the goal is to run as fast as possible in the vertical Y direction (as opposed to horizontal X direction in Ant) direction, and Ant-Energy, where the goal is to preserve energy. We detail how to construct these tasks in Appendix E. We use the same pretrained models for all these tasks. During adaptation, the model conditions on the $\mathcal{D}_{\text{few-shot}}$ and y^* for each task for optimization.

We evaluate ExPT on these tasks in the **poorest** setting. Table 3 shows that ExPT performs well on all tasks, where the median and mean scores are better than the best value in $\mathcal{D}_{\text{few-shot}}$, and the max score is close to 1. For the Ant, Ant- v_y , and Ant-Energy tasks, we visualize the behavior of the optimal designs that are discovered at <https://imgur.com/a/zpgI4YL>. When subject to the same policy-network (optimizing for horizontal X speed), the robots optimized for different objectives behave differently; the optimal Ant is capable of leaping forward to move quickly in X ; Ant- v_y is able to jump up to maximize speed in Y ; Ant-Energy is capable of ‘sitting down’ to conserve energy.

Table 3: ExPT’s performance on different objectives in D’Kitty and Ant domains. We pretrain one model for all tasks in the same domain. The performance is averaged across 3 seeds.

Task		D’Kitty	D’Kitty-45	D’Kitty-60	Ant	Ant- v_y	Ant-Energy
$\mathcal{D}_{\text{few-shot}}$ (best)		0.307	0.297	0.344	0.124	0.210	0.189
Median	ExPT	0.922 ± 0.009	0.611 ± 0.007	0.569 ± 0.010	0.686 ± 0.090	0.613 ± 0.009	0.635 ± 0.028
Max	ExPT	0.976 ± 0.004	0.954 ± 0.008	0.973 ± 0.004	0.965 ± 0.004	0.923 ± 0.049	0.950 ± 0.033
Mean	ExPT	0.871 ± 0.018	0.619 ± 0.016	0.584 ± 0.008	0.646 ± 0.061	0.599 ± 0.005	0.608 ± 0.025

B.5 Forward modeling versus Inverse modeling

As we mentioned in Section 2.2, two possible approaches exist to pretrain ExPT on synthetic data. We take the inverse modeling approach for ExPT throughout the paper, as we train ExPT to directly produce design inputs x ’s. In this section, we empirically validate our design choices by comparing ExPT with TNP-ED, its forward counterpart. TNP-ED’s

architecture is similar to ExPT’s in Figure 1, except that the target points now contain $x_{m+1:N}$ instead of $y_{m+1:N}$, the decoder is a 1-layer MLP, and the predicted outputs are $\hat{y}_{m+1:N}$. We call this model TNP-ED because a model named TNP (Nguyen and Grover, 2022) with a similar architecture was previously proposed in the context of meta-learning. We pretrain TNP-ED using a simple mean-squared error loss $\mathcal{L} = \sum_{i=m+1}^N (\hat{y}_i - y_i)^2$. After pretraining, we condition TNP-ED on $\mathcal{D}_{\text{few-shot}}$ and the best inputs in this dataset, and perform gradient ascent with respect to these inputs to obtain better points.

Table 4: Comparison of inverse modeling (ExPT) versus forward modeling (TNP-ED) on Ant and D’Kitty in **random** (left) and **poorest** (right) settings. We average the performance across 3 seeds.

	Baseline	D’Kitty	Ant		Baseline	D’Kitty	Ant
	$\mathcal{D}_{\text{few-shot}}(\text{best})$	0.883	0.563		$\mathcal{D}_{\text{few-shot}}(\text{best})$	0.307	0.124
Median	ExPT	0.902 ± 0.006	0.705 ± 0.018	Median	ExPT	0.922 ± 0.009	0.686 ± 0.090
	TNP-ED	0.770 ± 0.009	0.438 ± 0.007		TNP-ED	0.309 ± 0.000	0.197 ± 0.005
Mean	ExPT	0.865 ± 0.016	0.639 ± 0.026	Mean	ExPT	0.871 ± 0.018	0.646 ± 0.061
	TNP-ED	0.670 ± 0.037	0.451 ± 0.018		TNP-ED	0.405 ± 0.004	0.237 ± 0.005

Table 4 compares the performance of ExPT and TNP-ED on D’Kitty and Ant with respect to the median score and mean score. ExPT achieves significantly better performance in all metrics, especially in the **poorest** setting. This is because forward models suffer from poor out-of-distribution generalization, and performing gradient ascent on this model may result in points that have high values under the model but are very poor when evaluated using the true functions. This validates our inverse modeling approach.

Appendix C. ExPT pretraining details

Data generation For each domain, we pretrain ExPT for 10,000 iterations with 128 synthetic functions in each iteration, corresponding to a total number of 1,280,000 synthetic functions. For each function, we randomly sample 228 input x ’s from the unlabeled dataset $\mathcal{D}_{\text{unlabeled}}$ and generate the values y ’s from a Gaussian Process with an RBF kernel. To increase the diversity of synthetic data, we randomize the two hyperparameters, length scale $\ell \sim \mathcal{U}[5.0, 10.0]$ and function scale $\sigma \sim \mathcal{U}[1.0, 1.0]$, when generating each function. Additionally, we add Gaussian noises $\epsilon \sim \mathcal{N}(0, 0.1)$ to each input x sampled from $\mathcal{D}_{\text{unlabeled}}$ to enlarge the pretraining inputs. For each generated function, we use 100 points as context points and the remaining 128 as target points, and train the model to optimize (??). During the adaptation phase, we condition the pretrained ExPT model on the labeled few-shot dataset $\mathcal{D}_{\text{few-shot}}$ and the target function value y^* to generate designs x ’s.

Architectural details In all experiments, we use the same ExPT architecture. Before feeding to the Transformer encoder, we embed the (y, x) context pairs with a 1-layer MLP and embed the target y ’s with another 1-layer MLP. The transformer encoder has 4 layers with a hidden dimension of 128, 4 attention heads, GELU activation, and a dropout rate of 0.1. For the VAE model, we use a standard isotropic Gaussian distribution as the prior. Both the VAE encoder and VAE decoder have 4 layers with a hidden dimension of 512, and the latent variable z has a dimension of 32.

Optimization details In Design-Bench experiments, we pretrain ExPT for 10,000 iterations with 128 synthetic functions in each iteration. We use AdamW optimizer (Kingma and Ba, 2014; Loshchilov and Hutter) with a learning rate of $5e - 4$ and $(\beta_1, \beta_2) = (0.9, 0.99)$. We use a linear warmup schedule for 1000 steps, followed by a cosine-annealing schedule for 9000 steps.

Appendix D. Excluded Design-Bench tasks

D.1 Superconductor

We found the approximate oracle provided by Design-Bench not accurate enough to provide a reliable comparison of optimization methods on this task. Figure 4 plots the score values in the dataset against the score values predicted by the approximate oracle, which shows a weak correlation between these two values.




Figure 4: The correlation between the score values in the dataset (x -axis) and the score values predicted by the approximate oracle (y -axis) in Superconductor.

D.2 Hopper

As noted in previous works that use Design-Bench Krishnamoorthy et al. (2022), the oracle provided for the Hopper task is inconsistent with the true-dataset values. The outputs of the oracle on the dataset are skewed heavily towards low-function values, which makes it an unreliable task for evaluation.

D.3 ChEMBL

As observed in previous works (Trabucco et al., 2022; Krishnamoorthy et al., 2022), all methods produced nearly the same results on the ChEMBL task, so we excluded it in our experiments.

Appendix E. Additional ablation and analysis

E.1 Effects of GP hyperparameters

We empirically examine the impact of two GP hyperparameters, the variance σ and the length scale ℓ , on the performance of ExPT. Specifically, we evaluate the performance of ExPT on D’Kitty and Ant when σ is too small (ExPT-small- σ) or too large (ExPT-large- σ), and when ℓ is too small (ExPT-small- ℓ) or too large (ExPT-large- ℓ). In ExPT-small- σ and ExPT-large- σ , we sample σ from $\mathcal{U}[0.01, 0.1]$ and $\mathcal{U}[100, 200]$, respectively. In ExPT-small- ℓ and ExPT-large- ℓ , we sample ℓ from $\mathcal{U}[0.1, 1.0]$ and $\mathcal{U}[100, 200]$, respectively.

Table 5: Impact of σ and ℓ on ExPT performance on Ant and D’Kitty in **random** (left) and **poorest** (right) settings. We average the performance across 3 seeds.

	Baseline	D’Kitty	Ant		Baseline	D’Kitty	Ant
	$\mathcal{D}_{\text{few-shot}}(\text{best})$	0.883	0.563		$\mathcal{D}_{\text{few-shot}}(\text{best})$	0.307	0.124
Median	ExPT	0.902 ± 0.006	0.705 ± 0.018	Median	ExPT	0.922 ± 0.009	0.686 ± 0.090
	ExPT-small- σ	0.915 ± 0.006	0.661 ± 0.111		ExPT-small- σ	0.862 ± 0.064	0.656 ± 0.098
	ExPT-large- σ	0.797 ± 0.000	0.471 ± 0.012		ExPT-large- σ	0.792 ± 0.004	0.489 ± 0.019
	ExPT-small- ℓ	0.793 ± 0.004	0.459 ± 0.005		ExPT-small- ℓ	0.792 ± 0.006	0.462 ± 0.004
	ExPT-large- ℓ	0.795 ± 0.003	0.460 ± 0.003		ExPT-large- ℓ	0.795 ± 0.003	0.460 ± 0.004
Mean	ExPT	0.865 ± 0.016	0.639 ± 0.026	Mean	ExPT	0.871 ± 0.018	0.646 ± 0.061
	ExPT-small- σ	0.896 ± 0.016	0.630 ± 0.089		ExPT-small- σ	0.755 ± 0.085	0.606 ± 0.077
	ExPT-large- σ	0.752 ± 0.013	0.534 ± 0.015		ExPT-large- σ	0.726 ± 0.016	0.547 ± 0.012
	ExPT-small- ℓ	0.726 ± 0.018	0.518 ± 0.018		ExPT-small- ℓ	0.725 ± 0.019	0.529 ± 0.014
	ExPT-large- ℓ	0.725 ± 0.016	0.528 ± 0.006		ExPT-large- ℓ	0.722 ± 0.014	0.530 ± 0.011

The results in Table 5 show that overall, suboptimal values of σ and ℓ lead to a substantial drop in the performance of ExPT on both tasks. It is also noticeable that ℓ has a more significant influence on the performance than σ . In other words, the shape of the synthetic functions has a more critical impact on downstream performances than the magnitudes of the function values. A too small ℓ or large ℓ results in synthetic functions that exhibit either excessive oscillations or excessive smoothness, leading to poor generalization to downstream functions.

E.2 Construction of new D’Kitty and Ant tasks

This section details how we constructed new objectives from the original D’Kitty and Ant that we used to evaluate ExPT in Section B.4. For each new objective, we apply the corresponding oracle to the inputs x ’s in the original dataset to create the dataset for the objective.

Ant tasks In Ant, the original goal is to design a morphology that allows the Ant robot to run as fast as possible in the x (horizontal) direction. The objective function is the sum of rewards in 100 time steps, where the reward R at each time step is defined as:

$$R = \text{Forward reward} + \text{Survival reward} - \text{Control cost} - \text{Contact cost}, \quad (1)$$

where Forward reward = $(x_t - x_{t-1})/dt$ is the velocity of the Ant in the x direction.

In Ant- v_y , the reward at each time step is similar, except that Forward reward = $(y_t - y_{t-1})/dt$ is the velocity of the Ant in the y (vertical) direction. In other words, we aim to design morphologies that allow the robot to run fast in the y direction.

In Ant-Energy, the reward at each time step is:

$$R = 1 + \text{Survival reward} - \text{Control cost} - \text{Contact cost}, \quad (2)$$

which means we incentivize the robot to conserve energy instead of running fast.

D’Kitty tasks In D’Kitty, the goal is to design a morphology that allows the D’Kitty robot to reach a fixed target location, and the objective function f is the Euclidean distance to the target. In the original D’Kitty task, the target location is on the vertical line from the starting point. In the two new tasks D’Kitty-45 and D’Kitty-60, the target locations are 45 deg and 60 deg away from the original target, respectively.




Figure 5: Different D’Kitty tasks. The red dot denotes the starting location, and the green dots are the target locations.

E.3 ExPT with sequential sampling

A significant advantage of ExPT is its ability to adapt to any objective function purely through in-context learning. This means that the model can refine its understanding of the underlying objective function given more data points in a very efficient manner. In this section, we explore an alternative optimization scheme for ExPT, namely *sequential sampling*, which explicitly utilizes the in-context learning ability of the model. Specifically, instead of producing $Q = 256$ inputs simultaneously, we sample one by one sequentially. That is, we condition the model on $\mathcal{D}_{\text{few-shot}}$ and y^* to sample the first point, evaluate the point using the black-box function, and add the point together with its score to the context set. We repeat this process for 256 times.

Table 6 shows that ExPT with sequential sampling performs better than simultaneous sampling on D’Kitty and Ant in both **random** and **poor** settings. Especially on Ant in the **poorest** setting, ExPT-Sequential achieves improvements of 20% and 19% over ExPT in terms of the median and mean performance, respectively. Intuitively, as we add more data points to the context set, ExPT-Sequential is able to update its understanding of the structure of the objective function, consequently leading to improved performance.

Table 6: Comparison of simultaneous (ExPT) and sequential (ExPT-Seq) sampling on Ant and D’Kitty in **random** (left) and **poorest** (right) settings. We average the performance across 3 seeds.

	Baseline	D’Kitty	Ant		Baseline	D’Kitty	Ant
	$\mathcal{D}_{\text{few-shot}}(\text{best})$	0.883	0.563		$\mathcal{D}_{\text{few-shot}}(\text{best})$	0.307	0.124
Median	ExPT	0.902 ± 0.006	0.705 ± 0.018	Median	ExPT	0.922 ± 0.009	0.686 ± 0.090
	ExPT-Seq	0.903 ± 0.005	0.719 ± 0.013		ExPT-Seq	0.928 ± 0.012	0.822 ± 0.055
Mean	ExPT	0.865 ± 0.016	0.639 ± 0.026	Mean	ExPT	0.871 ± 0.018	0.646 ± 0.061
	ExPT-Seq	0.872 ± 0.010	0.669 ± 0.017		ExPT-Seq	0.923 ± 0.011	0.767 ± 0.048

E.4 Effects of $|\mathcal{D}_{\text{unlabeled}}|$

We empirically examine the effects of the size of $\mathcal{D}_{\text{unlabeled}}$ on the downstream performance of ExPT. Specifically, we subsample the x ’s in the public dataset with a ratio $r \in \{0.01, 0.1, 0.2, 0.5, 1.0\}$. Adaptation and evaluation are the same as in Section 4.

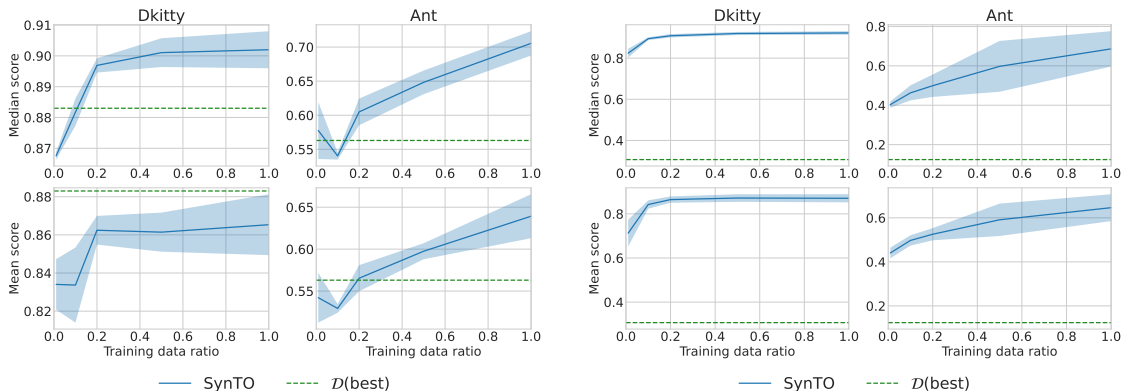


Figure 6: The performance of ExPT on Dkitty and Ant in the **random** (left) and **poorest** (right) setting when we vary the training data ration r . We average the performance across 3 seeds.

Figure 6 shows the median and mean performance of ExPT on Dkitty and Ant in both **random** and **poorest** settings with respect to the ratio r . In the **random** setting, ExPT is able to reach or surpass the best data point in the few-shot dataset by using as few as 0.2 of the pretraining data. In the **poorest** setting, ExPT performs better than the best dataset point with only 0.01 of the pretraining data. Moreover, the performance improves consistently as the pretraining data size increases, suggesting that we can achieve even better performance by simply using more unlabeled data for pretraining. This result highlights the unique capability of ExPT of learning from unlabeled data, providing new opportunities for solving challenging optimization problems where unlabeled data is plentiful but labeled data is scarce.

E.5 Sorting context and target points

In the main experiments in Section 4, for each generated function during pretraining, we sample 228 points that we divide randomly into 100 context points and 128 target points. However, at adaptation, we condition on target output values that are likely to be higher than the best input value in the context set. Therefore, it is natural to sort the context points and target points during pretraining, so that the target inputs always have higher values than the context inputs. We denote this pretraining mechanism as ExPT-sorted.

Table 7: Comparison of pretraining on randomly divided context and target points (ExPT) versus sorted context and target points (ExPT-sorted) on Ant and D’Kitty in **random** (left) and **poorest** (right) settings. We average the performance across 3 seeds.

		Baseline	D’Kitty	Ant			Baseline	D’Kitty	Ant
		$\mathcal{D}_{\text{few-shot}}(\text{best})$	0.883	0.563			$\mathcal{D}_{\text{few-shot}}(\text{best})$	0.307	0.124
Median	ExPT	0.902 ± 0.006	0.705 ± 0.018		Median	ExPT	0.922 ± 0.009	0.686 ± 0.090	
	ExPT-Sorted	0.811 ± 0.019	0.631 ± 0.015			ExPT-Sorted	0.911 ± 0.003	0.685 ± 0.044	
Mean	ExPT	0.865 ± 0.016	0.639 ± 0.026		Mean	ExPT	0.871 ± 0.018	0.646 ± 0.061	
	ExPT-Sorted	0.794 ± 0.020	0.590 ± 0.014			ExPT-Sorted	0.900 ± 0.003	0.642 ± 0.035	

Table 7 shows that ExPT-sorted underperforms ExPT in the **random** setting, while performing very similarly in the **poorest** setting. This indicates that learning to predict any points provides a better and more general pretraining objective than only learning to predict points with high values.

E.6 Comparisons with more baselines

In addition to the baselines in Section 4, we compare ExPT with 3 variants of Gradient Ascent, a method that was considered in previous works (Trabucco et al., 2022; Kumar and Levine, 2020; Trabucco et al., 2021; Krishnamoorthy et al., 2022). The Grad. Asc baseline simply learns a forward model and finds an optimal x^* by taking 200 gradient-ascent steps to improve an existing input x . The two variants Grad. Min and Grad. Mean create ensembles of forward models and perform gradient ascent using the min and mean ensemble predictions.

Tables 8 and 9 show the performance of ExPT and all baselines in the **random** and **poorest** settings. We see that while the gradient ascent methods perform well on certain tasks, with good performance on the TF-Bind8 task in particular, ExPT is still the best performing method in all settings and metrics.

Appendix F. Compute

All training is done on 10 AMD EPYC 7313 CPU cores and one NVIDIA RTX A5000 GPU.

Appendix G. Reproducibility

We made a strong effort to ensure that our work can be reproduced properly. In Section 3, we provide a comprehensive description of our methodology, while in Section 4 and Appendix B, we provide the specifics of our pretraining and evaluation setup, as well as our choice of hyperparameters. We compare our approach with various baseline methods from different

Table 8: Comparison of ExPT and the baselines on the few-shot **random** setting of 4 Design-Bench tasks. We report median, max, and mean performance across 3 random seeds. Higher scores are better. **Blue** denotes the best entry in the column, and **Violet** denotes the second best.

	Baseline	D’Kitty	Ant	TF Bind 8	TF Bind 10	Mean score (\uparrow)
	$\mathcal{D}_{\text{few-shot}}(\text{best})$	0.883	0.563	0.439	0.466	—
Median	MINs	0.859 \pm 0.014	0.485 \pm 0.152	0.416 \pm 0.019	0.468 \pm 0.014	0.557 \pm 0.050
	COMs	0.752 \pm 0.007	0.411 \pm 0.012	0.371 \pm 0.001	0.468 \pm 0.000	0.501 \pm 0.005
	BONET	0.852 \pm 0.013	0.597 \pm 0.119	0.441 \pm 0.003	0.483 \pm 0.009	0.593 \pm 0.036
	BDI	0.592 \pm 0.020	0.396 \pm 0.018	0.540 \pm 0.032	0.438 \pm 0.034	0.492 \pm 0.026
	GP-qEI	0.842 \pm 0.058	0.550 \pm 0.007	0.439 \pm 0.000	0.467 \pm 0.000	0.575 \pm 0.016
	Grad. Asc	0.403 \pm 0.134	0.088 \pm 0.017	0.492 \pm 0.017	0.492 \pm 0.018	0.369 \pm 0.0465
	Grad. Min	0.712 \pm 0.028	0.220 \pm 0.035	0.504 \pm 0.025	0.465 \pm 0.008	0.475 \pm 0.024
	Grad. Mean	0.437 \pm 0.180	0.150 \pm 0.037	0.551 \pm 0.029	0.485 \pm 0.018	0.406 \pm 0.066
	ExPT	0.902 \pm 0.006	0.705 \pm 0.018	0.473 \pm 0.014	0.477 \pm 0.014	0.639 \pm 0.013
Max	MINs	0.930 \pm 0.010	0.890 \pm 0.017	0.814 \pm 0.030	0.639 \pm 0.017	0.818 \pm 0.019
	COMs	0.920 \pm 0.010	0.841 \pm 0.044	0.686 \pm 0.152	0.656 \pm 0.020	0.776 \pm 0.057
	BONET	0.909 \pm 0.012	0.888 \pm 0.024	0.887 \pm 0.053	0.702 \pm 0.006	0.847 \pm 0.024
	BDI	0.918 \pm 0.006	0.806 \pm 0.094	0.906 \pm 0.074	0.532 \pm 0.023	0.791 \pm 0.049
	GP-qEI	0.896 \pm 0.000	0.887 \pm 0.000	0.513 \pm 0.104	0.647 \pm 0.011	0.736 \pm 0.029
	Grad. Asc	0.775 \pm 0.032	0.240 \pm 0.032	0.923 \pm 0.005	0.675 \pm 0.017	0.653 \pm 0.0215
	Grad. Min	0.822 \pm 0.053	0.434 \pm 0.092	0.960 \pm 0.002	0.632 \pm 0.009	0.712 \pm 0.039
	Grad. Mean	0.829 \pm 0.009	0.337 \pm 0.063	0.957 \pm 0.010	0.668 \pm 0.034	0.698 \pm 0.029
	ExPT	0.973 \pm 0.005	0.970 \pm 0.004	0.933 \pm 0.036	0.677 \pm 0.048	0.888 \pm 0.023
Mean	MINs	0.624 \pm 0.025	0.009 \pm 0.013	0.415 \pm 0.030	0.465 \pm 0.015	0.378 \pm 0.021
	COMs	0.515 \pm 0.050	0.020 \pm 0.006	0.369 \pm 0.003	0.471 \pm 0.004	0.344 \pm 0.016
	BONET	0.837 \pm 0.023	0.579 \pm 0.024	0.448 \pm 0.011	0.484 \pm 0.009	0.587 \pm 0.017
	BDI	0.570 \pm 0.032	0.385 \pm 0.012	0.536 \pm 0.032	0.444 \pm 0.027	0.484 \pm 0.026
	GP-qEI	0.505 \pm 0.006	0.019 \pm 0.001	0.439 \pm 0.001	0.473 \pm 0.002	0.359 \pm 0.003
	Grad. Asc	0.400 \pm 0.073	0.090 \pm 0.018	0.513 \pm 0.014	0.492 \pm 0.017	0.374 \pm 0.031
	Grad. Min	0.599 \pm 0.068	0.221 \pm 0.034	0.531 \pm 0.015	0.462 \pm 0.009	0.453 \pm 0.032
	Grad. Mean	0.527 \pm 0.079	0.150 \pm 0.038	0.569 \pm 0.028	0.438 \pm 0.017	0.421 \pm 0.041
	ExPT	0.865 \pm 0.016	0.639 \pm 0.026	0.476 \pm 0.010	0.474 \pm 0.015	0.614 \pm 0.017

approaches on multiple tasks in Design-Bench (Trabucco et al., 2022) with distinct properties. Our results are averaged over 3 seeds and we also report the standard deviation. Additionally, we conduct several ablation experiments to examine how sensitive ExPT is to different hyperparameters.

Appendix H. Broader impact

The field of offline black-box optimization can have positive impacts in many spheres, including in drug-discovery, nuclear reactor design, and optimal robot design. The few-shot setting that we introduce in this work is also highly relevant to these fields which have large quantities of unlabelled data, but only a limited quantity of labelled data points. It is also worth noting however, that it is possible to use black-box optimization in general for malicious purposes such as to produce chemicals with harmful properties. Even though our work does not directly enable such use cases, this possibility should be taken into account when applying ExPT and similar frameworks to these kinds of impactful real-world scenarios.

Table 9: Comparison of ExPT and the baselines on the few-shot **poorest** setting of 4 Design-Bench tasks. We report the median, max, and mean performance across 3 random seeds. Higher scores are better. **Blue** denotes the best entry in the column, and **Violet** denotes the second best.

	Baseline	D’Kitty	Ant	TF Bind 8	TF Bind 10	Mean score (\uparrow)
	$\mathcal{D}_{\text{few-shot}}(\text{best})$	0.307	0.124	0.124	0.000	—
Median	MINs	0.480 \pm 0.156	0.316 \pm 0.040	0.437 \pm 0.007	0.463 \pm 0.003	0.424 \pm 0.052
	COMs	0.733 \pm 0.023	0.401 \pm 0.026	0.111 \pm 0.000	0.459 \pm 0.006	0.426 \pm 0.014
	BONET	0.310 \pm 0.000	0.236 \pm 0.047	0.319 \pm 0.018	0.461 \pm 0.017*	0.332 \pm 0.021
	BDI	0.309 \pm 0.000	0.192 \pm 0.012	0.365 \pm 0.000	0.454 \pm 0.017	0.330 \pm 0.007
	GP-qEI	0.883 \pm 0.000	0.565 \pm 0.001	0.439 \pm 0.000	0.467 \pm 0.000	0.589 \pm 0.000
	Grad. Asc	0.741 \pm 0.026	0.321 \pm 0.012	0.425 \pm 0.064	0.419 \pm 0.073	0.477 \pm 0.044
	Grad. Min	0.806 \pm 0.004	0.454 \pm 0.061	0.357 \pm 0.040	0.376 \pm 0.079	0.498 \pm 0.046
	Grad. Mean	0.742 \pm 0.054	0.472 \pm 0.066	0.350 \pm 0.014	0.395 \pm 0.019	0.489 \pm 0.038
	ExPT	0.922 \pm 0.009	0.686 \pm 0.090	0.552 \pm 0.042	0.489 \pm 0.013	0.662 \pm 0.039
Max	MINs	0.841 \pm 0.014	0.721 \pm 0.031	0.962 \pm 0.019	0.648 \pm 0.025	0.793 \pm 0.022
	COMs	0.931 \pm 0.022	0.843 \pm 0.020	0.124 \pm 0.000	0.739 \pm 0.057	0.659 \pm 0.025
	BONET	0.929 \pm 0.031	0.557 \pm 0.118	0.809 \pm 0.038	0.519 \pm 0.039*	0.704 \pm 0.057
	BDI	0.939 \pm 0.002	0.693 \pm 0.109	0.913 \pm 0.000	0.596 \pm 0.020	0.785 \pm 0.033
	GP-qEI	0.896 \pm 0.000	0.887 \pm 0.000	0.439 \pm 0.000	0.645 \pm 0.021	0.717 \pm 0.005
	Grad. Asc	0.837 \pm 0.038	0.684 \pm 0.071	0.821 \pm 0.077	0.568 \pm 0.019	0.728 \pm 0.052
	Grad. Min	0.910 \pm 0.009	0.801 \pm 0.029	0.842 \pm 0.066	0.555 \pm 0.028	0.777 \pm 0.033
	Grad. Mean	0.882 \pm 0.028	0.807 \pm 0.046	0.747 \pm 0.055	0.542 \pm 0.039	0.745 \pm 0.042
	ExPT	0.946 \pm 0.018	0.965 \pm 0.004	0.873 \pm 0.035	0.615 \pm 0.022	0.850 \pm 0.020
Mean	MINs	0.623 \pm 0.051	0.015 \pm 0.017	0.464 \pm 0.009	0.463 \pm 0.002	0.391 \pm 0.020
	COMs	0.607 \pm 0.021	0.033 \pm 0.003	0.109 \pm 0.001	0.454 \pm 0.004	0.301 \pm 0.007
	BONET	0.490 \pm 0.023	0.234 \pm 0.052	0.318 \pm 0.018	0.459 \pm 0.018	0.375 \pm 0.028
	BDI	0.364 \pm 0.004	0.215 \pm 0.021	0.369 \pm 0.000	0.453 \pm 0.018	0.350 \pm 0.011
	GP-qEI	0.533 \pm 0.001	0.018 \pm 0.000	0.439 \pm 0.000	0.470 \pm 0.002	0.365 \pm 0.001
	Grad. Asc	0.659 \pm 0.069	0.334 \pm 0.018	0.432 \pm 0.061	0.427 \pm 0.042	0.463 \pm 0.048
	Grad. Min	0.794 \pm 0.003	0.454 \pm 0.051	0.374 \pm 0.018	0.386 \pm 0.044	0.502 \pm 0.029
	Grad. Mean	0.702 \pm 0.083	0.467 \pm 0.050	0.356 \pm 0.023	0.405 \pm 0.018	0.483 \pm 0.044
	ExPT	0.871 \pm 0.018	0.646 \pm 0.061	0.549 \pm 0.032	0.488 \pm 0.011	0.639 \pm 0.031