Generative Adversarial Model-Based Optimization via Source Critic Regularization

Michael S. Yao, Yimeng Zeng, Hamsa Bastani,
Jacob Gardner, James Gee, and Osbert Bastani

{myao2199, obastani}@seas.upenn.edu

University of Pennsylvania, Philadephia, PA 19104

NEURAL INFORMATION PROCESSING SYSTEMS

Challenges of Offline Generative Design

In real-world tasks like molecule and robotic design, optimizing objective functions can be costly or impractical. To address this, recent work explores **offline policy optimization**, leveraging prior observations to find the best designs without ever having access to the objective function. However, offline methods often struggle with the **distribution shift** between the distribution of the design-generating policy and the offline observations. To overcome this, we introduce a **generative adversarial model-based optimization (GAMBO)**: task-agnostic approach to reliably optimize against an offline surrogate model for generative tasks.

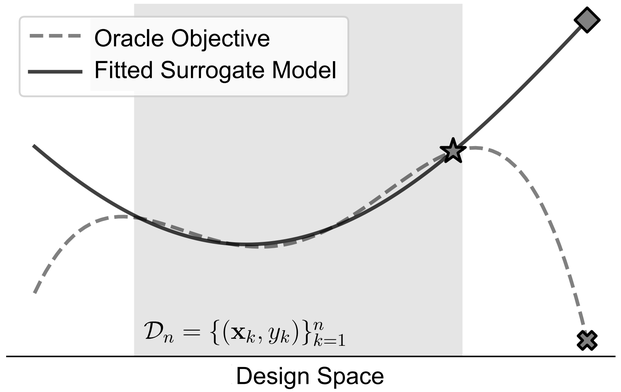

Problem Formulation

Model Based Optimization

We consider the class of problems of the form

$$\mathbf{x}^* = \underset{\mathbf{x}\in\mathcal{X}}{\operatorname{argmax}}\, f(\mathbf{x})$$

In settings where the **oracle function** f is expensive to evaluate, we can instead learn an **offline surrogate model** f_θ from a dataset of prior observations $\mathcal{D}_n = \{(\mathbf{x}_k, y_k)\}_{k=1}^n$ and instead solve the related problem

$$\mathbf{x}^* = \underset{\mathbf{x}\in\mathcal{X}}{\operatorname{argmax}}\, f_\theta(\mathbf{x})$$

How "reliable" is the surrogate model f_θ over $\mathcal{X}$?

The **Wasserstein Distance** is a measure of similarity between probability distributions commonly used in generative models. We use the Wasserstein distance as a proxy for how similar generated designs are to previously observed designs.

Our Approach: Constrained Optimization

We formulate a constrained MBO problem according to

$$\begin{aligned} \text{minimize}_{\boldsymbol{x}\in\mathcal{X}} \quad & -f_\theta(\mathbf{x}) \\ \text{subject to} \quad & \mathbb{E}_{\mathbf{x}'\in\mathcal{D}_n}[c^*(\mathbf{x}')] - c^*(\mathbf{x}) \le 0 \end{aligned}$$

where $c^*(x)\colon \mathcal{X} \to \mathbb{R}$ is a **source critic function** used to approximate the Wasserstein distance between the generated and offline designs. In other words, we seek to maximize f_θ over the feasible search space of designs that are **at least as in-distribution** as the average design in the offline dataset $\mathcal{D}_n$.

We use an **augmented Lagrangian approach** from standard convex optimization theory. That is, our modified objective is

$$x^* = \underset{\mathbf{x}\in\mathcal{X}}{\operatorname{argmax}}\, \mathcal{L}(\mathbf{x};\lambda)$$

where $\mathcal{L}(\mathbf{x};\lambda)$ is the Lagrangian

$$\mathcal{L}(\mathbf{x};\lambda) \coloneqq f_\theta(x) - \lambda\left[\mathbb{E}_{\mathbf{x}'\in\mathcal{D}_n}[c^*(\mathbf{x}')] - c^*(\mathbf{x})\right]$$

The parameter λ balances the tradeoff between maximizing against f_θ and staying in-distribution with respect to $\mathcal{D}_n$. Our algorithm **GAMBO** automatically calculates the optimal value of λ!

In theory, GAMBO can be leveraged with any offline optimizer. We leverage GAMBO for **Bayesian optimization** (i.e., GABO) and **gradient ascent** (i.e., GAGA) in our experiments.

Experimental Evaluation

We evaluate GABO and GAGA against state-of-the-art offline optimization algorithms on the following generative design tasks:

1. **Branin**: Design points to maximize the Branin function.
2. **Molecule**: Design molecules with maximal penalized LogP score (a metric of molecule hydrophobicity).
3. **TF-Bind-8**: Design an 8-bp DNA sequence maximizing the binding efficiency with a particular transcription factor.
4. **GFP**: Design proteins with maximal green fluorescence.
5. **UTR**: Design a 5'UTR sequence that maximizes the gene expression of the encoded protein.
6. **ChEMBL**: Design a molecule with maximal predicted bioactivity according to a specific assay.
7. **D'Kitty**: Design a D'Kitty robot for optimal navigation.
8. **Warfarin**: Design the optimal warfarin dose for a patient based on clinical and pharmacogenetic variables.

Method	Branin	LogP	TF-Bind-8*	GFP*	UTR*	ChEMBL*	D'Kitty	Warfarin	Rank
$\mathcal{D}$ (best)	-13.0	11.3	0.439	3.53	7.12	0.61	0.88	-0.19 ± 1.96	—
Grad.	-245.1 ± 81.3	-5.37 ± 1.44	0.429 ± 0.023	3.18 ± 0.88	6.82 ± 0.21	-1.95 ± 0.00	0.57 ± 0.19	0.86 ± 1.09	9.0
L-BFGS	-29.6 ± 0.0	3.82 ± 32.6	0.527 ± 0.140	3.51 ± 0.70	6.48 ± 1.20	-1.95 ± 0.00	0.31 ± 0.00	0.73 ± 1.83	8.5
CMA-ES	-8.6 ± 3.6	5.04 ± 6.83	0.438 ± 0.131	1.43 ± 0.00	6.39 ± 0.11	-1.95 ± 0.00	0.31 ± 0.00	-25.0 ± 150	10.6
Anneal	-9.6 ± 1.5	8.76 ± 0.15	0.807 ± 0.094	3.64 ± 0.03	5.01 ± 0.31	-1.95 ± 0.00	0.55 ± 0.18	**0.91 ± 0.08**	6.8
BO	-11.0 ± 7.8	-52.5 ± 88.8	0.586 ± 0.193	1.43 ± 0.00	5.65 ± 1.30	0.59 ± 0.10	0.61 ± 0.15	0.16 ± 1.67	8.8
TuRBO	-21.0 ± 5.1	-45.1 ± 93.8	0.564 ± 0.194	1.43 ± 0.00	6.53 ± 1.19	**0.65 ± 0.00**	0.44 ± 0.18	0.05 ± 0.11	9.0
BONET	-26.1 ± 0.9	10.8 ± 0.33	0.282 ± 0.000	**3.74 ± 0.00**	**9.12 ± 0.07**	0.55 ± 0.13	0.78 ± 0.00	—	5.7
DDOM	-6677 ± 6360	-4.23 ± 1.28	0.460 ± 0.030	1.43 ± 0.00	5.56 ± 0.02	0.54 ± 0.15	0.51 ± 0.20	-0.32 ± 0.40	11.1
COM	-3099 ± 32.6	**30.8 ± 19.5**	0.439 ± 0.000	3.62 ± 0.00	6.65 ± 0.43	0.63 ± 0.01	**0.90 ± 0.02**	0.72 ± 0.97	5.5
RoMA	-32.7 ± 18.4	6.37 ± 1.39	0.433 ± 0.040	3.37 ± 0.27	6.66 ± 0.98	0.50 ± 0.14	0.30 ± 0.27	-0.70 ± 0.02	9.4
BDI	-1050 ± 0.0	-0.20 ± 0.00	0.311 ± 0.000	3.26 ± 0.82	5.61 ± 0.00	0.48 ± 0.00	0.67 ± 0.00	-24.8 ± 233	10.8
ExPT	-57.2 ± 38.6	-15.9 ± 24.1	0.571 ± 0.076	1.43 ± 0.00	6.77 ± 1.38	0.56 ± 0.06	0.66 ± 0.20	-34.6 ± 61.4	9.1
BootGen	—	-13.0 ± 15.1	**0.942 ± 0.022**	3.10 ± 0.73	8.30 ± 0.93	0.59 ± 0.07	—	—	6.2
ROMO	-2614 ± 739.9	-20.5 ± 19.2	0.382 ± 0.203	3.55 ± 0.13	5.73 ± 1.42	**0.65 ± 0.00**	0.64 ± 0.27	-0.71 ± 2.10	9.6
GAGA	-2.9 ± 2.2	-68.6 ± 109.8	0.571 ± 0.120	**3.74 ± 0.00**	5.89 ± 1.42	-1.95 ± 0.00	0.89 ± 0.00	0.01 ± 0.14	7.4
GABO	**-2.6 ± 1.1**	21.3 ± 33.2	0.570 ± 0.131	3.60 ± 0.40	7.51 ± 0.39	0.60 ± 0.07	0.71 ± 0.01	0.60 ± 1.80	**3.8**

Table 1: One-Shot Oracle Evaluation. We show one-shot design scores (higher is better) on 8 real-world offline optimization tasks. Average rank (lower is better) shown in the right-most common.

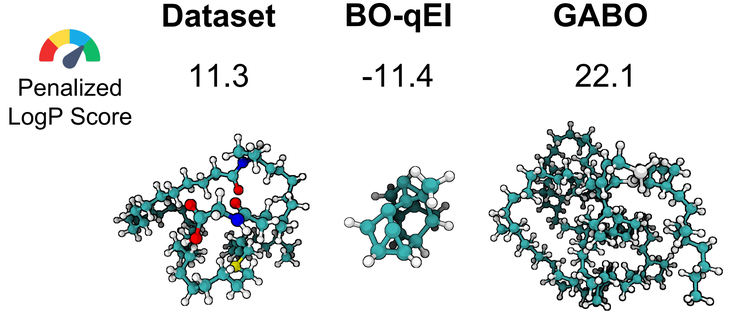

Figure 1: Example Proposed Molecules. GABO designs molecules that outperform previously observed molecules and molecules performed by other optimization algorithms.

Discussion and Conclusion

GAMBO dynamically adjusts the regularization strength according to the optimization trajectory. When designs are in-distribution, GAMBO relaxes the constraint to explore more of the design space. When designs start to look "wacky," GAMBO tightens the constraint to stay more "in-distribution."

What's next?

1. Can we generalize our dual optimization approach?
2. Can GAMBO be used for other tasks, like RHLF and LLM preference optimization?

Penn Engineering University of Pennsylvania
Penn Medicine
NSF
NIH

December 2024

Generative Adversarial Model-Based Optimization via Source Critic Regularization

Michael S. Yao, Yimeng Zeng, Hamsa Bastani,

Jacob R. Gardner, James C. Gee, Osbert Bastani

{myao2199, obastani}@seas.upenn.edu

NEURAL INFORMATION PROCESSING SYSTEMS

Generative Design

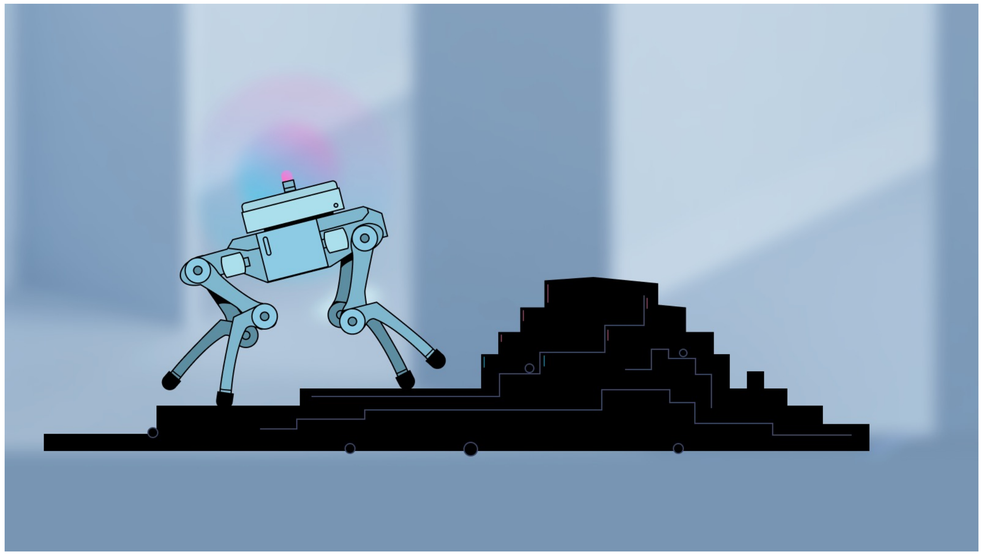

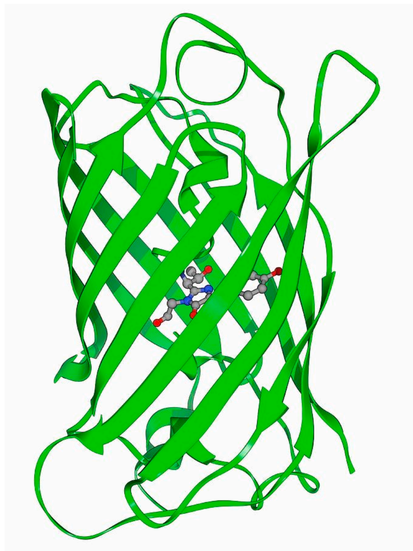

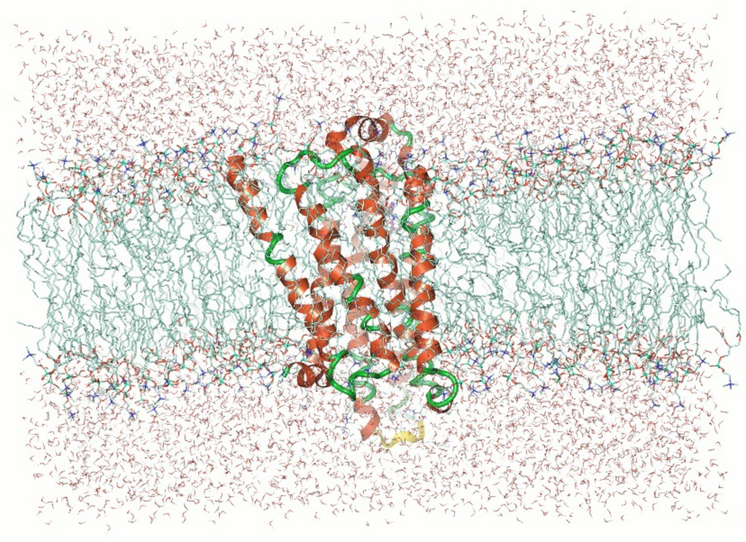

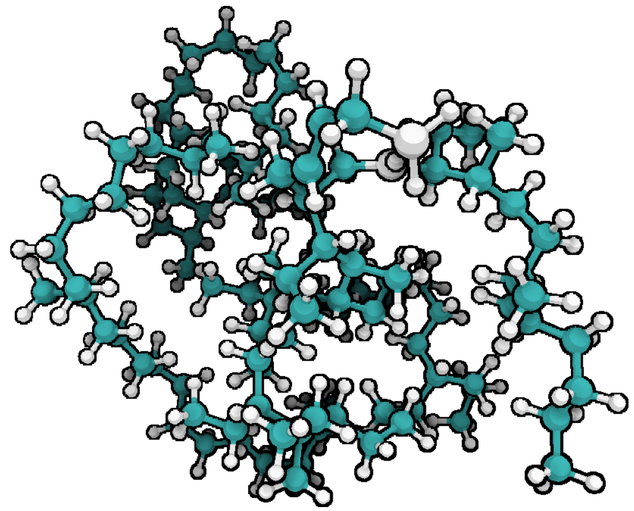

NEURAL INFORMATION PROCESSING SYSTEMS

Simulations and surrogates aren’t perfect!

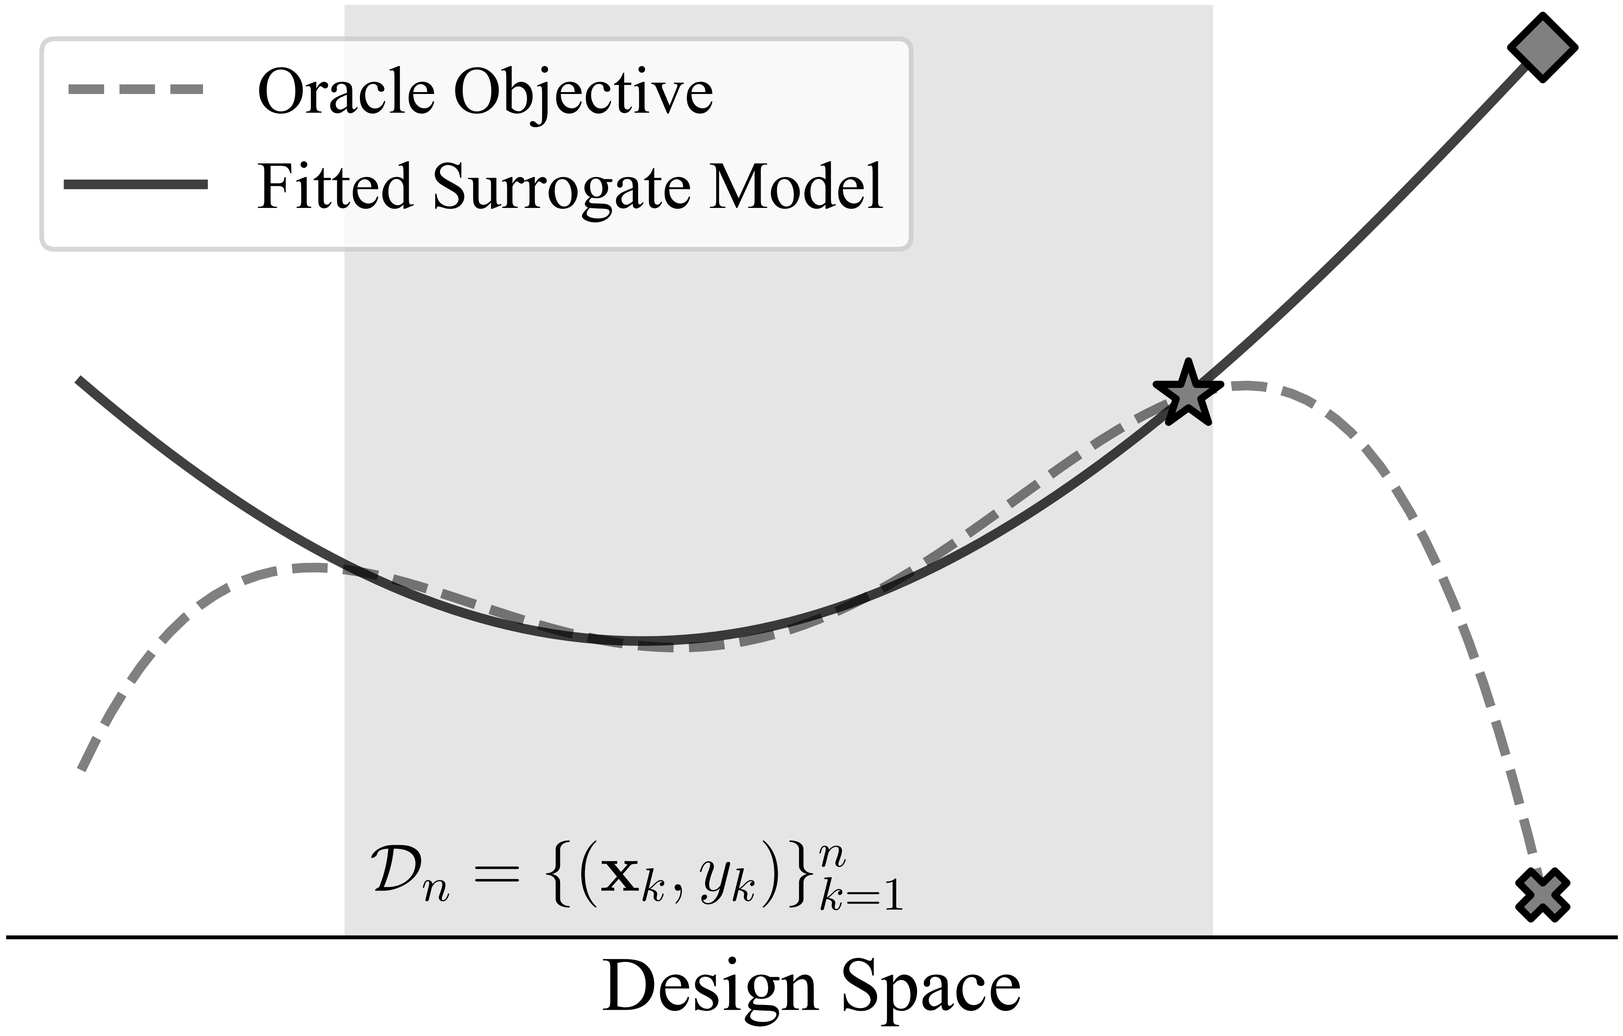

NEURAL INFORMATION PROCESSING SYSTEMS

How can we measure offline distribution shift?

$$W_1(p_n, q_n) = \inf_{\sigma} \frac{1}{n} \sum_{i=1}^{n} \left\| z'_{\sigma(i)} - z_i \right\| = \sup_{\|c\|_L \leq K} \mathbb{E}_{z'_j \sim p_n} \left[c(z'_j) \right] - \mathbb{E}_{z_i \sim q_n} \left[c(z_i) \right]$$

$$p_n = \frac{1}{n} \sum_{j=1}^{n} \delta(z - z'_j)$$

Reference Distribution

$$q_n = \frac{1}{n} \sum_{i=1}^{n} \delta(z - z_i)$$

Generated Distribution

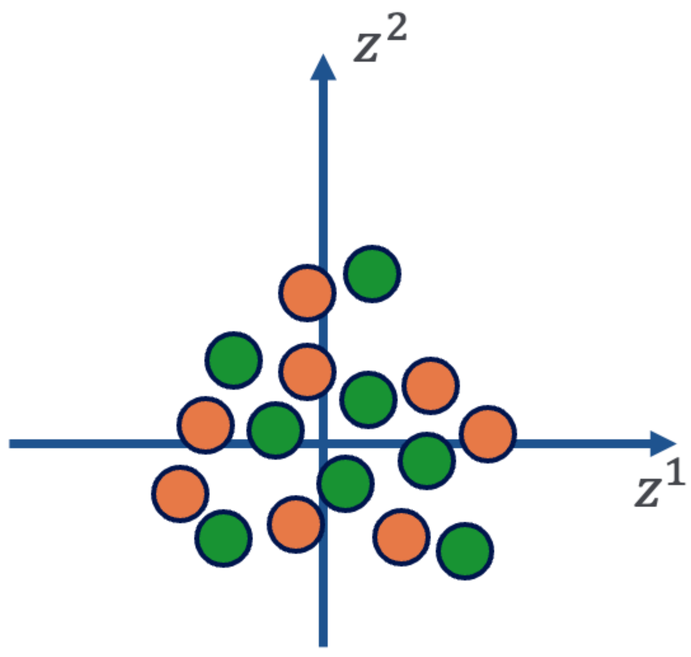

$W_1(p_n, q_n)$ small

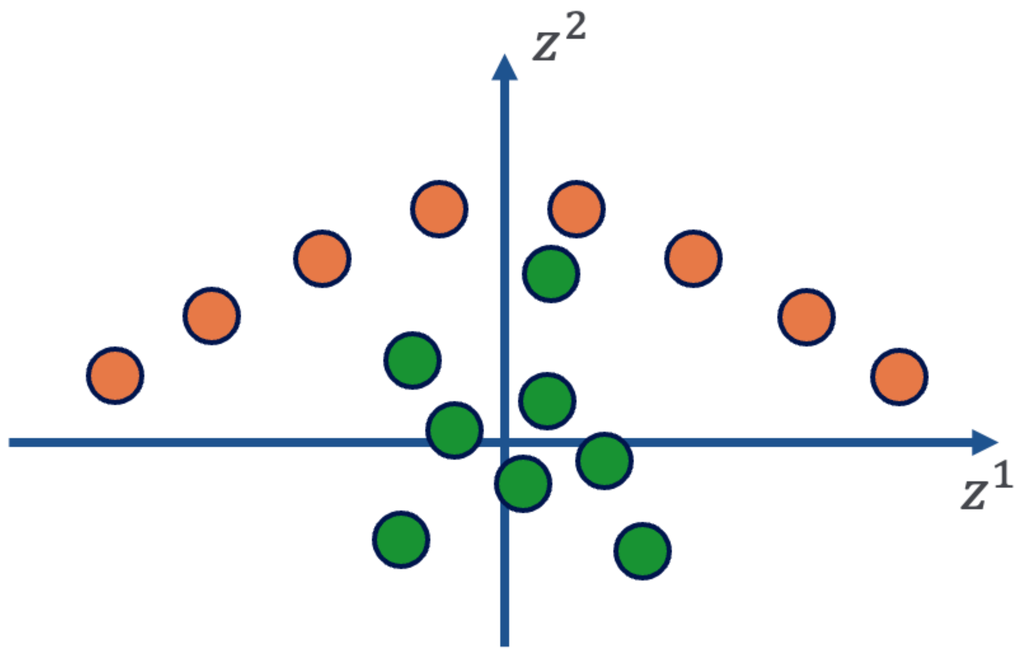

$W_1(p_n, q_n)$ large

NEURAL INFORMATION PROCESSING SYSTEMS

How can we measure offline distribution shift?

$$\begin{aligned} \text{minimize}_{z \in \mathcal{Z}} \quad & -f_\theta(z) \\ \text{subject to} \quad & \mathbb{E}_{z' \in P}[c^*(z')] - c^*(z) \leq 0 \end{aligned}$$

$-f_\theta(z)$

$\mathbb{E}_{z' \in P}[c^*(z')] - c^*(z)$

How do we balance this tradeoff?

Propose a design that maximizes the surrogate

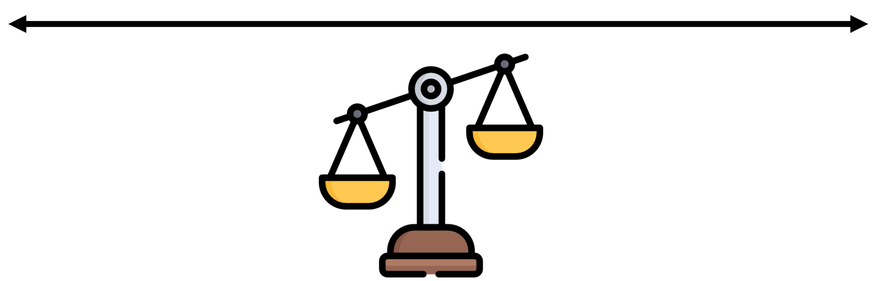

Propose a design that "looks like" other designs

Generative Adversarial Model-Based Optimization (GAMBO)

NEURAL INFORMATION PROCESSING SYSTEMS

How can we measure offline distribution shift?

Method	Branin	LogP	TF-Bind-8*	GFP*	UTR*	ChEMBL*	D'Kitty	Warfarin	Rank
$\mathcal{D}$ (best)	-13.0	11.3	0.439	3.53	7.12	0.61	0.88	-0.19 ± 1.96	—
Grad.	-245.1 ± 81.3	-5.37 ± 1.44	0.429 ± 0.023	3.18 ± 0.88	6.82 ± 0.21	-1.95 ± 0.00	0.57 ± 0.19	0.86 ± 1.09	9.0
L-BFGS	-29.6 ± 0.0	3.82 ± 32.6	0.527 ± 0.140	3.51 ± 0.70	6.48 ± 1.20	-1.95 ± 0.00	0.31 ± 0.00	0.73 ± 1.83	8.5
CMA-ES	-8.6 ± 3.6	5.04 ± 6.83	0.438 ± 0.131	1.43 ± 0.00	6.39 ± 0.11	-1.95 ± 0.00	0.31 ± 0.00	-25.0 ± 150	10.6
Anneal	-9.6 ± 1.5	8.76 ± 0.15	0.807 ± 0.094	3.64 ± 0.03	5.01 ± 0.31	-1.95 ± 0.00	0.55 ± 0.18	**0.91** ± **0.08**	6.8
BO	-11.0 ± 7.8	-52.5 ± 88.8	0.586 ± 0.193	1.43 ± 0.00	5.65 ± 1.30	0.59 ± 0.10	0.61 ± 0.15	0.16 ± 1.67	8.8
TuRBO	-21.0 ± 5.1	-45.1 ± 93.8	0.564 ± 0.194	1.43 ± 0.00	6.53 ± 1.19	**0.65** ± **0.00**	0.44 ± 0.18	0.05 ± 0.11	9.0
BONET	-26.1 ± 0.9	10.8 ± 0.33	0.282 ± 0.000	**3.74** ± **0.00**	**9.12** ± **0.07**	0.55 ± 0.13	0.78 ± 0.00	—	5.7
DDOM	-6677 ± 6360	-4.23 ± 1.28	0.460 ± 0.030	1.43 ± 0.00	5.56 ± 0.02	0.54 ± 0.15	0.51 ± 0.20	-0.32 ± 0.40	11.1
COM	-3099 ± 32.6	**30.8** ± **19.5**	0.439 ± 0.000	3.62 ± 0.00	6.65 ± 0.43	0.63 ± 0.01	**0.90** ± **0.02**	0.72 ± 0.97	5.5
RoMA	-32.7 ± 18.4	6.37 ± 1.39	0.433 ± 0.040	3.37 ± 0.27	6.66 ± 0.98	0.50 ± 0.14	0.30 ± 0.27	-0.70 ± 0.02	9.4
BDI	-1050 ± 0.0	-0.20 ± 0.00	0.311 ± 0.000	3.26 ± 0.82	5.61 ± 0.00	0.48 ± 0.00	0.67 ± 0.00	-24.8 ± 233	10.8
ExPT	-57.2 ± 38.6	-15.9 ± 24.1	0.571 ± 0.076	1.43 ± 0.00	6.77 ± 1.38	0.56 ± 0.06	0.66 ± 0.20	-34.6 ± 61.4	9.1
BootGen	—	-13.0 ± 15.1	**0.942** ± **0.022**	3.10 ± 0.73	8.30 ± 0.93	0.59 ± 0.07	—	—	6.2
ROMO	-2614 ± 739.9	-20.5 ± 19.2	0.382 ± 0.203	3.55 ± 0.13	5.73 ± 1.42	**0.65** ± **0.00**	0.64 ± 0.27	-0.71 ± 2.10	9.6
GAGA	-2.9 ± 2.2	-68.6 ± 109.8	0.571 ± 0.120	**3.74** ± **0.00**	5.89 ± 1.42	-1.95 ± 0.00	0.89 ± 0.00	0.01 ± 0.14	7.4
GABO	**-2.6** ± **1.1**	21.3 ± 33.2	0.570 ± 0.131	3.60 ± 0.40	7.51 ± 0.39	0.60 ± 0.07	0.71 ± 0.01	0.60 ± 1.80	**3.8**

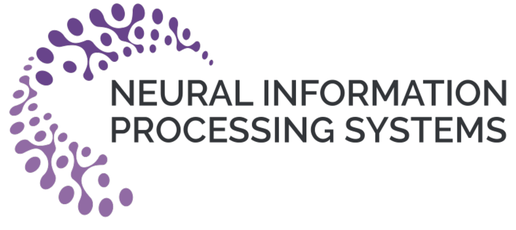

Thank you!

Michael Yao
myao2199@seas.upenn.edu

Yimeng Zeng

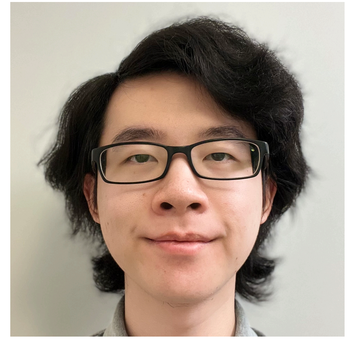

Hamsa Bastani

Osbert Bastani
obastani@seas.upenn.edu

Jacob Gardner

James Gee

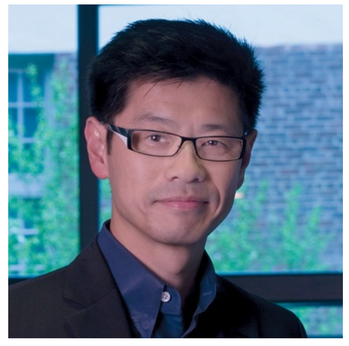

Penn Engineering
University of Pennsylvania

Penn Medicine

NSF

NIH

Poster Session 4 | Thursday 12/12 | 4:30 – 7:30 PST

NEURAL INFORMATION PROCESSING SYSTEMS