

Design Editing for Offline Model-based Optimization

Ye Yuan ^{*}[†], Youyuan Zhang ^{*}, Can (Sam) Chen, Haolun Wu,
Zixuan (Melody) Li, Jianmo Li, James J. Clark, Xue Liu





Problem Background

- Create objects that exhibits specific target properties.
 - For example: Develop a new superconductor material to achieve higher critical temperature.
- Evaluation can be expensive or dangerous, so assume access only to an offline dataset of designs and their property scores.
 - For example: some pairs of existing superconductor materials and their corresponding critical temperatures.
- **Offline Model-based Optimization (MBO):** find a design (superconductor material) to maximize its property (critical temperature) with the offline dataset only.



Problem Formulation

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}),$$

- where $f(\cdot)$ denotes the unknown objective function, and $\mathbf{x} \in \mathcal{X}$ denotes a candidate design.
- An offline dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ is available, where \mathbf{x}_i represents a specific design, such as superconductor material, and y_i represents the corresponding score, like critical temperature.



Acquirement of Pseudo Design Candidates

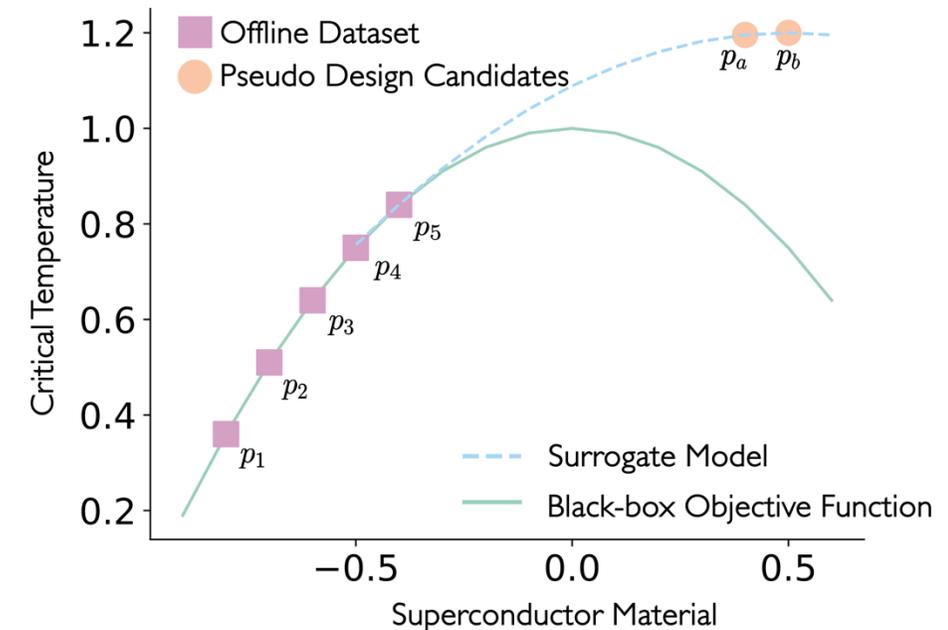
We employ the forward approach to acquire pseudo design candidates:

- 1) Fit a DNN surrogate $f_{\theta}(\cdot)$ to \mathcal{D} .

$$\theta^* = \arg \min_{\theta} \frac{1}{N} \sum_{i=1}^N (f_{\theta}(\mathbf{x}_i) - y_i)^2.$$

- 2) Perform gradient ascent:

$$\mathbf{x}_t = \mathbf{x}_{t-1} + \eta \nabla_{\mathbf{x}} f_{\theta}(\mathbf{x}) \Big|_{\mathbf{x}=\mathbf{x}_t}, \quad \text{for } t \in [1, T].$$



Out-of-distribution issue: The surrogate models often extrapolate unreliably when optimizing outside the training data distribution, resulting in erroneously high predicted property scores but low ground-truth scores



Motivation

- Can we explicitly edit pseudo design candidates, which are acquired from the naive forward approach, back into the valid distribution?
 - How do we model the valid distribution?
 - How do we effectively edit the pseudo design candidates?



Design Editing through Diffusion Prior

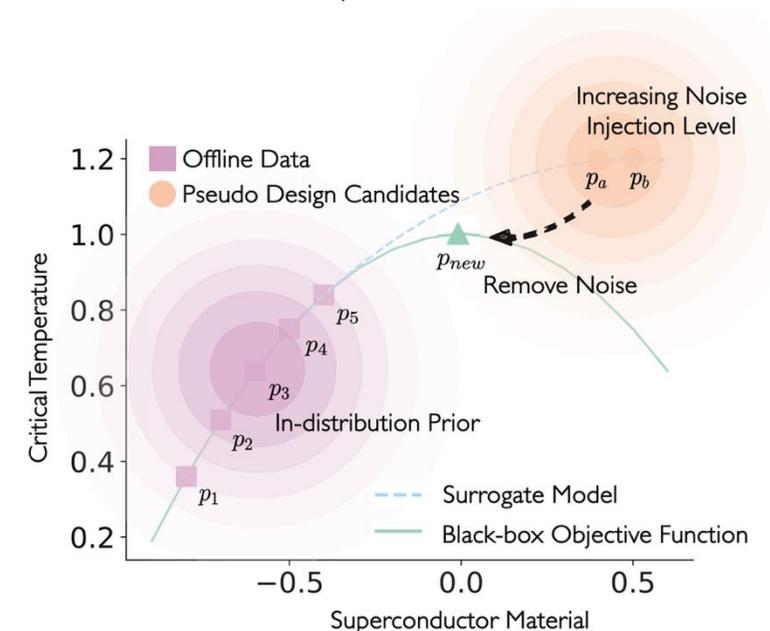
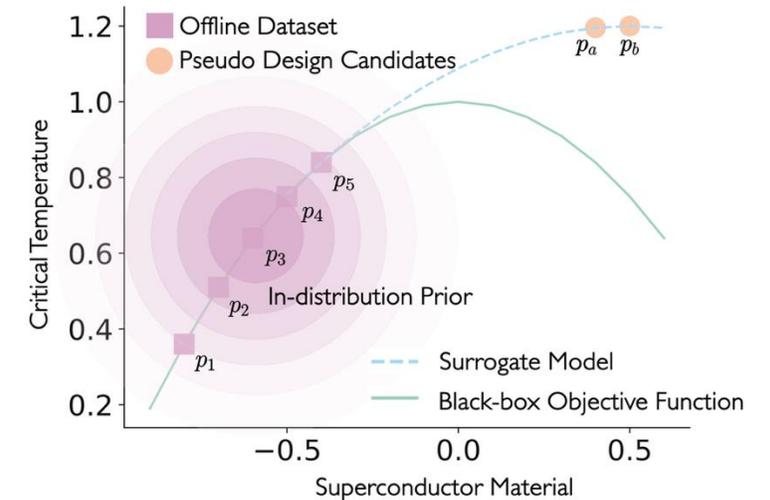
- A diffusion model is employed to capture the distribution of the offline dataset, functioning as an in-distribution prior:

$$\phi^* = \arg \min_{\phi} \mathbb{E}_t [\lambda(t) \mathbb{E}_{\mathbf{x}_0, y} [\mathbb{E}_{\mathbf{x}_t | \mathbf{x}_0} [\|\mathbf{s}_{\phi}(\mathbf{x}_t, t, y) - \nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t | \mathbf{x}_0)\|^2]]].$$

- To edit a pseudo design candidate $\mathbf{x}^{(p)}$, we perturb it by introducing noise at a specific time m :

$$\mathbf{x}_{\text{perturb}}^{(p)} = \sqrt{\bar{\alpha}_m} \mathbf{x}^{(p)} + \sqrt{1 - \bar{\alpha}_m} \epsilon.$$

- The perturbed design is used as the starting point. A final optimized design is synthesized by using numerical solver for the backward denoising process.





Experiment: Tasks

- **Superconductor:** develop a superconductor with 86 components to maximize the critical temperature.
- **Ant Morphology:** design a robot with 60 components to maximize the moving speed.
- **D’Kitty Morphology:** design a robot with 56 components to maximize the moving speed.
- **Levy:** optimize a 60-dimension continuous vector to maximize the inverse Levy black-box function.
- **TF Bind 8:** discover an 8-unit DNA sequence to maximize the binding affinity.
- **TF Bind 10:** find a 10-unit DNA sequence to maximize the binding affinity.
- **NAS:** find the optimal neural network architecture to enhance test accuracy.



Experiment: Evaluation Metrics

- Generate 256 designs for each approach and identify the design with the maximum score. Report the 100th percentile normalized ground-truth score:

$$y_n = \frac{y - y_{min}}{y_{max} - y_{min}},$$

where y_{min} and y_{max} represent the minimum and maximum scores within the entire unobserved dataset, respectively.

- Report the best design in the offline dataset for better comparison, denoted as $\mathcal{D}(\text{best})$.



Experimental Results: Continuous Tasks

Table 1: Experimental results on continuous tasks for comparison.

Method	Superconductor	Ant Morphology	D’Kitty Morphology	Levy
$\mathcal{D}(\text{best})$	0.399	0.565	0.884	0.613
BO-qEI	0.402 ± 0.034	0.819 ± 0.000	0.896 ± 0.000	0.810 ± 0.016
CMA-ES	0.465 ± 0.024	1.214 ± 0.732	0.724 ± 0.001	0.887 ± 0.025
REINFORCE	0.481 ± 0.013	0.266 ± 0.032	0.562 ± 0.196	0.564 ± 0.090
Mean	0.505 ± 0.013	0.940 ± 0.014	0.956 ± 0.014	0.984 ± 0.023
Min	0.501 ± 0.019	0.918 ± 0.034	0.942 ± 0.009	0.964 ± 0.023
COMs	0.481 ± 0.028	0.842 ± 0.037	0.926 ± 0.019	0.936 ± 0.025
ROMA	0.509 ± 0.015	0.916 ± 0.030	0.929 ± 0.013	0.976 ± 0.019
NEMO	0.502 ± 0.002	0.955 ± 0.006	0.952 ± 0.004	0.969 ± 0.019
BDI	0.513 ± 0.000	0.906 ± 0.000	0.919 ± 0.000	0.938 ± 0.000
IOM	0.518 ± 0.020	0.922 ± 0.030	0.944 ± 0.012	0.988 ± 0.021
ICT	0.503 ± 0.017	0.961 ± 0.007	0.968 ± 0.020	0.879 ± 0.018
Tri-mentoring	0.514 ± 0.018	0.948 ± 0.014	0.966 ± 0.010	0.924 ± 0.035
PGS	0.563 ± 0.058	0.949 ± 0.017	0.966 ± 0.013	0.963 ± 0.027
CbAS	0.503 ± 0.069	0.876 ± 0.031	0.892 ± 0.008	0.938 ± 0.037
Auto CbAS	0.421 ± 0.045	0.882 ± 0.045	0.906 ± 0.006	0.797 ± 0.033
MIN	0.499 ± 0.017	0.445 ± 0.080	0.892 ± 0.011	0.761 ± 0.037
DDOM	0.486 ± 0.013	0.952 ± 0.007	0.941 ± 0.006	0.927 ± 0.031
BONET	0.437 ± 0.022	0.976 ± 0.012	0.954 ± 0.012	0.918 ± 0.025
DEMO_(ours)	0.525 ± 0.009	0.968 ± 0.009	0.970 ± 0.007	1.007 ± 0.015

Our method achieves the competitive performance on four continuous tasks. Each method is run for 8 independent trials. The mean and standard deviation are reported.



Experimental Results: Discrete Tasks and Rankings

Table 2: Experimental results on discrete tasks, and ranking on all tasks for comparison.

Method	TF Bind 8	TF Bind 10	NAS	Rank Mean	Rank Median
$\mathcal{D}(\text{best})$	0.439	0.467	0.436		
BO-qEI	0.798 ± 0.083	0.652 ± 0.038	1.079 ± 0.059	13.9/19	16/19
CMA-ES	0.953 ± 0.022	0.670 ± 0.023	0.985 ± 0.079	9.1/19	7/19
REINFORCE	0.948 ± 0.028	0.663 ± 0.034	-1.895 ± 0.000	15.1/19	19/19
Mean	0.895 ± 0.020	0.654 ± 0.028	0.663 ± 0.058	9.3/19	9/19
Min	0.931 ± 0.036	0.634 ± 0.033	0.708 ± 0.027	10.7/19	11/19
COMs	0.474 ± 0.053	0.625 ± 0.010	0.796 ± 0.029	13.1/19	14/19
ROMA	0.921 ± 0.040	0.669 ± 0.035	0.934 ± 0.025	7.9/19	7/19
NEMO	0.942 ± 0.003	0.708 ± 0.010	0.735 ± 0.012	6.7/19	7/19
BDI	0.870 ± 0.000	0.605 ± 0.000	0.722 ± 0.000	12.1/19	13/19
IOM	0.870 ± 0.074	0.648 ± 0.025	0.411 ± 0.044	10.0/19	10/19
ICT	0.958 ± 0.008	0.691 ± 0.023	0.667 ± 0.091	7.7/19	6/19
Tri-mentoring	0.970 ± 0.001	0.722 ± 0.017	0.759 ± 0.102	5.4/19	4/19
PGS	0.981 ± 0.015	0.658 ± 0.021	0.727 ± 0.033	5.4/19	7/19
CbAS	0.927 ± 0.051	0.651 ± 0.060	0.683 ± 0.079	12.0/19	12/19
Auto CbAS	0.910 ± 0.044	0.630 ± 0.045	0.506 ± 0.074	15.6/19	16/19
MIN	0.905 ± 0.052	0.616 ± 0.021	0.717 ± 0.046	15.4/19	16/19
DDOM	0.961 ± 0.024	0.640 ± 0.029	0.737 ± 0.014	9.4/19	10/19
BONET	0.975 ± 0.004	0.681 ± 0.035	0.724 ± 0.008	8.0/19	6/19
DEMO_(ours)	0.982 ± 0.016	0.762 ± 0.058	0.753 ± 0.017	2.1/19	1/19

Our methods achieve the best result in two discrete tasks. We also achieve the best mean ranking and median ranking among all baselines.



Conclusion

- We propose Design Editing for Offline Model-based Optimization to trade-off between the overestimation caused by the extrapolation error of the surrogate model and the overconservatism by only generating new designs from the diffusion prior.
- We validate the design editing process can effectively leverage information from both the surrogate and the diffusion prior. Our approach demonstrates effectiveness in addressing MBO problem, yielding competitive results in multiple tasks.



Thanks for your attention!



Paper



Code



McGill
UNIVERSITY



Mila