ICLR

Pareto-Conditioned Diffusion Models for Offline Multi-Objective Optimization

Jatan Shrestha*, Santeri Heiskanen*, Kari Hepola,
Severi Rissanen, Pekka Jääskeläinen, Joni Pajarinen

Aalto University

Tampere University

Motivation

- Many real-world optimization problems are **multi-objective**
    - *Biological Sequence Design*
    - *Neural Architecture Search*
- We aim to find a set of all **non-dominated solutions** that form the **Pareto Front**
- **Hypervolume** measures the quality of a solution set with respect to the reference point

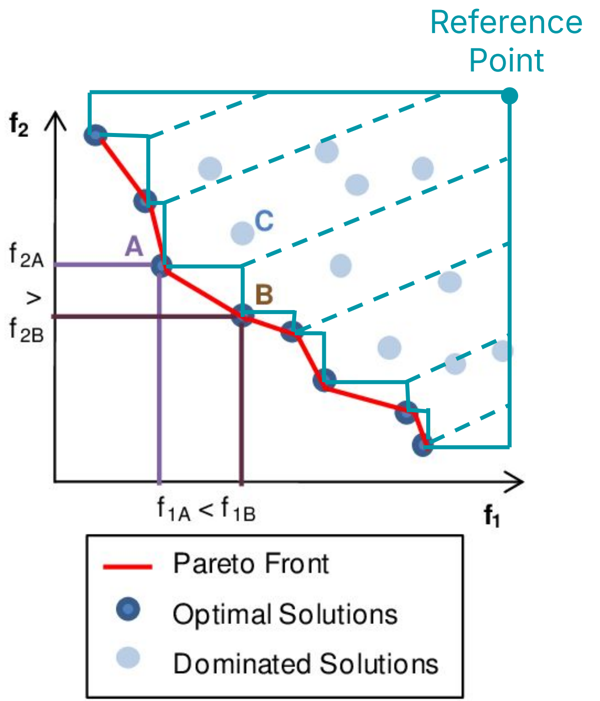

Offline Multi-Objective Optimization

- **Challenge:** Evaluating real-world designs is prohibitively expensive or risky
- **The "Offline" Setting:** Optimize using only a static, pre-collected dataset

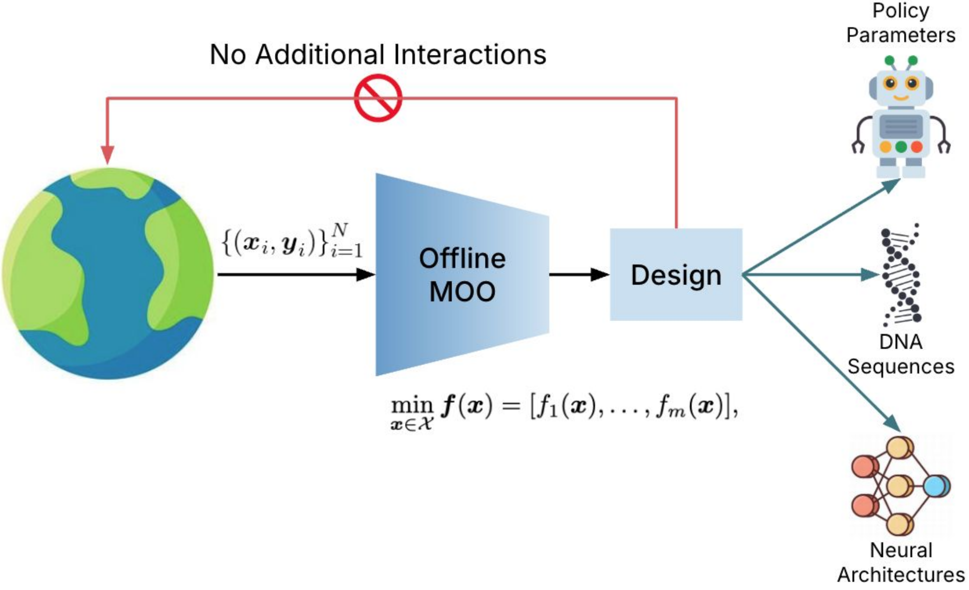

Prior Work

- **The Dominant Paradigm:** Multi-stage pipelines relying on explicit **surrogate models** or **predictors**
- **The Bottleneck:** Performance is limited by the accuracy of the surrogate model or predictor

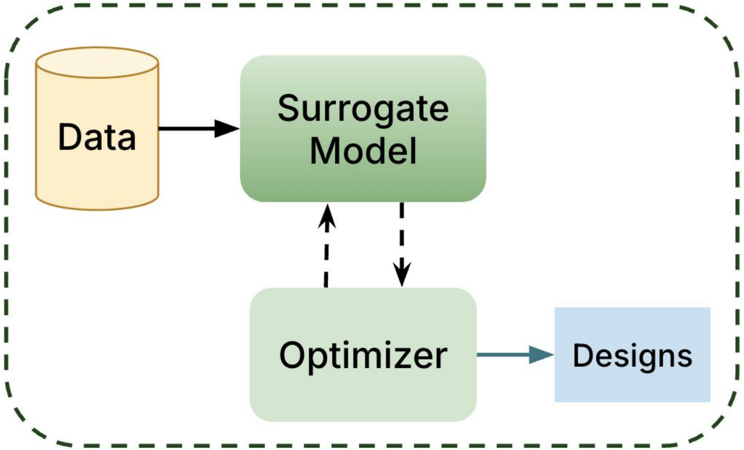

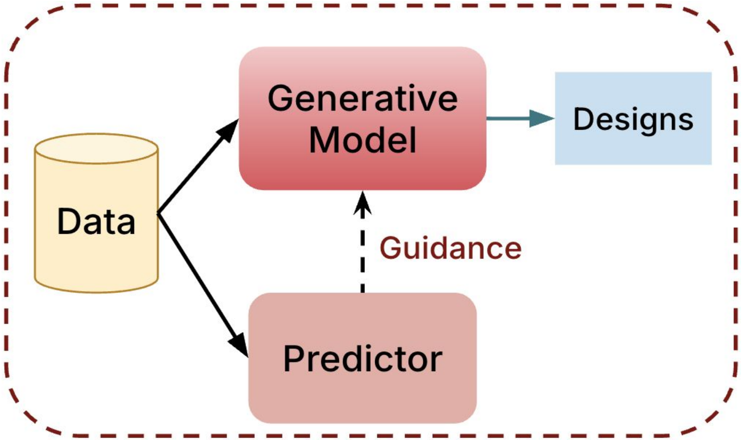

Pareto-Conditioned Diffusion Models

- **Our Solution:** We propose **Pareto-Conditioned Diffusion (PCD),** which reframes offline MOO as a **conditional sampling problem**

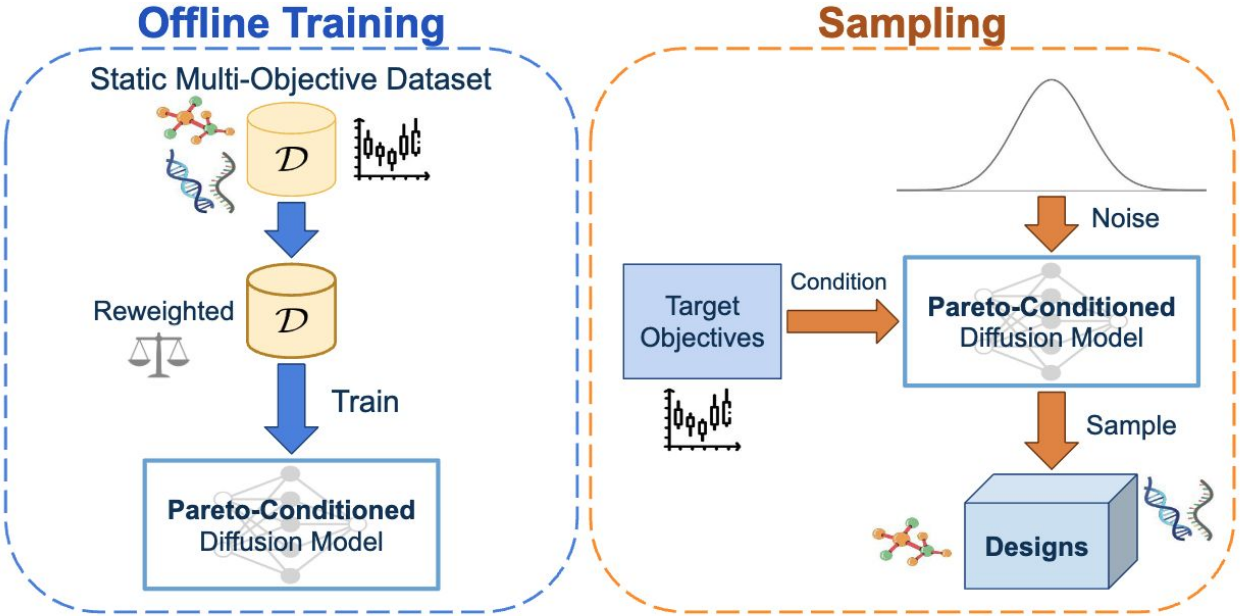

Training via Multi-Objective Reweighting

- **Goal:** Focus on **optimal** designs rather than average ones
- **Method:** Reweight the dataset using the **dominance number** of each point

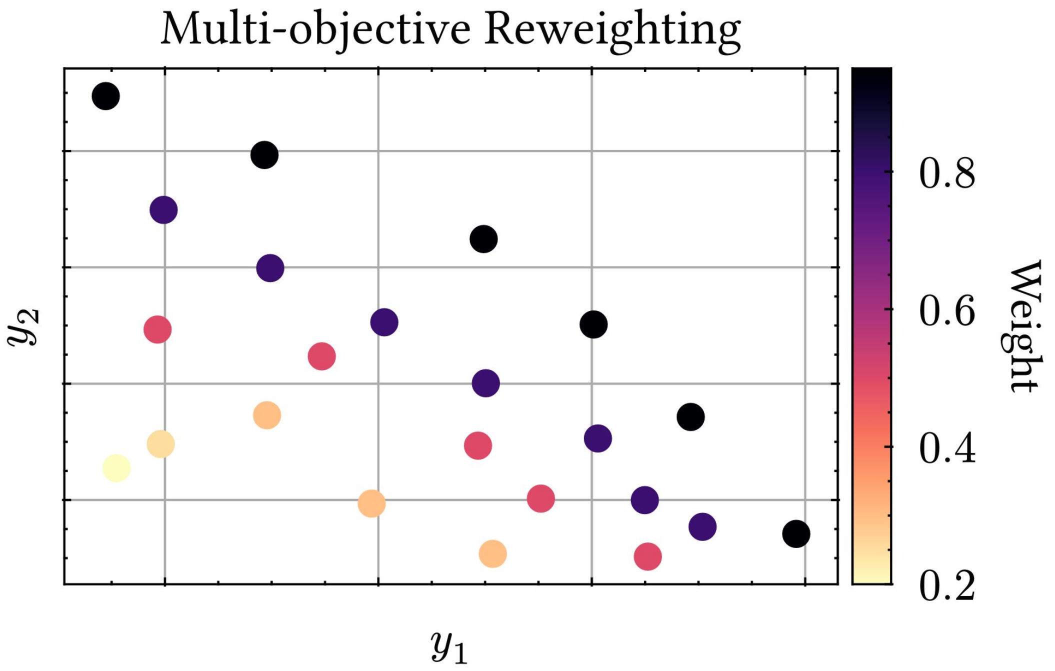

Generating Conditioning Points

1) Partition solution space via direction vectors

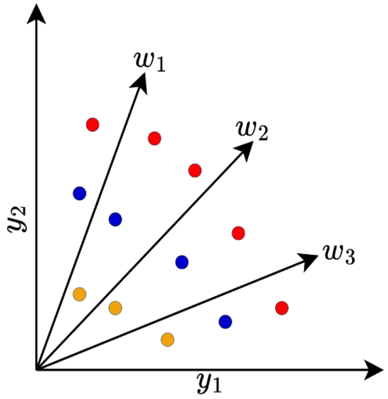

2) Assign points to direction vectors

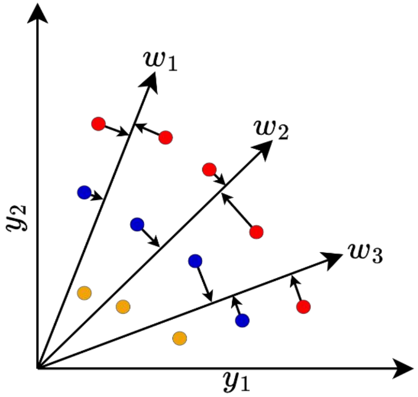

3) Extrapolate and add noise

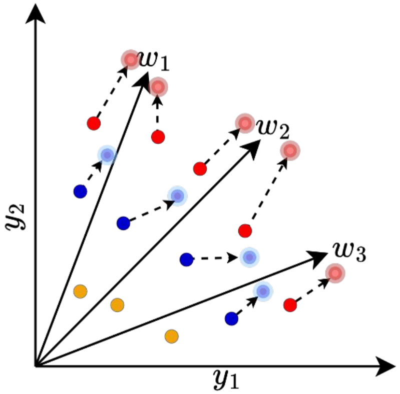

Sampling via Predictor-free Guidance

- Generate **designs** that are highly faithful to the **target trade-offs**

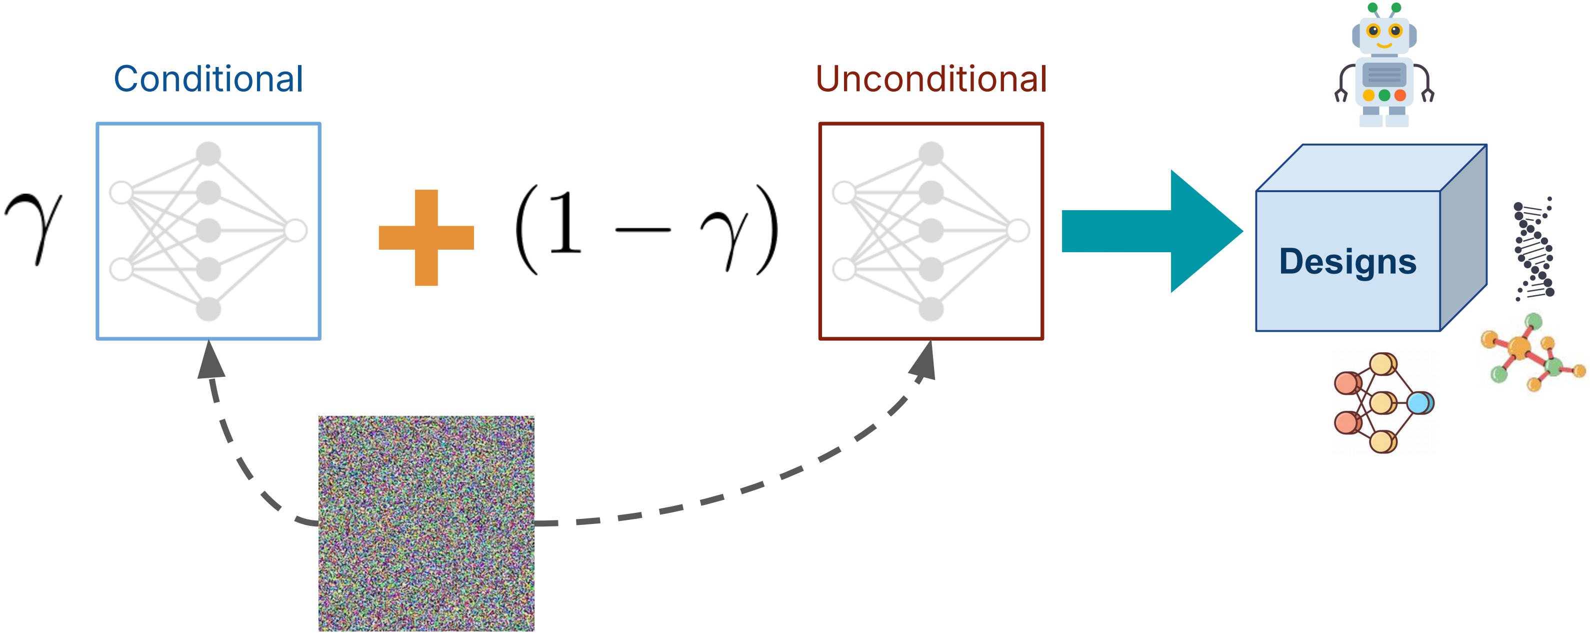

Main Results

- **PCD** is highly **competitive** and **consistent** across diverse offline MOO tasks with a single set of hyperparameters

Method	Synthetic	MORL	RE	Scientific	MONAS	Avg. rank
$\mathcal{D}$(best)	5.45±0.19	**1.70±0.27**	2.60±0.07	9.35±0.14	11.53±0.06	7.43±0.05
MOBO	8.69±0.30	14.60±0.42	10.00±0.33	6.75±0.47	8.11±0.80	8.81±0.34
E2E + GN	7.33±0.55	5.70±2.14	7.06±0.32	5.35±1.38	9.33±0.53	7.82±0.40
E2E + PC	5.93±0.25	<u>3.50±1.22</u>	6.22±0.33	4.30±1.32	6.60±0.40	6.01±0.29
E2E	6.16±0.30	9.70±2.08	6.06±0.30	4.20±1.40	<u>5.13±0.22</u>	<u>5.71±0.16</u>
MH + GN	8.82±0.53	8.90±2.16	8.14±0.94	5.05±2.14	12.57±0.40	9.84±0.33
MH + PC	8.87±0.45	10.90±1.08	6.74±0.68	6.15±0.91	7.46±0.30	7.68±0.33
MH	6.18±0.53	8.00±1.41	6.14±0.29	5.80±0.89	5.88±0.49	6.10±0.22
MM + COMs	8.02±0.47	3.60±1.29	6.54±0.17	**3.85±0.68**	7.22±0.43	6.80±0.13
MM + ICT	6.73±0.46	9.10±1.95	5.44±0.32	5.05±0.74	8.42±0.40	7.08±0.13
MM + IOM	5.16±0.51	12.70±0.91	5.76±0.52	4.40±1.15	5.77±0.50	5.80±0.20
MM + TM	6.55±0.82	7.90±2.16	5.78±0.25	5.90±1.29	7.87±0.39	6.91±0.20
MM	6.07±0.50	9.50±0.79	5.94±0.41	6.55±0.93	**4.97±0.46**	5.80±0.21
ParetoFlow	**2.44±0.28**	8.50±1.32	<u>1.74±0.17</u>	9.05±0.27	11.19±0.52	6.74±0.23
PCD (ours)	<u>3.38±0.20</u>	5.50±3.30	**1.51±0.13**	<u>4.05±0.33</u>	7.54±0.50	**4.80±0.30**

Conclusion

- We introduced **PCD**, reframing offline MOO as a **conditional sampling problem**
- Leverages **dominance-based reweighting** and **extrapolated targets** to generate diverse, novel solutions
- **PCD** achieves highly **competitive** and **consistent** performance across diverse benchmarks

Project website