# Conservative Pareto Set Amortization for Offline Multi-Objective Optimization

Ji Cheng · City University of Hong Kong

OffPSL learns from a fixed archive, then returns **one design for each requested trade-off**.

CANDIDATE GENERATION

**60 ms**

**to return 256 designs**

PGD-MOO takes **7.83 s** *on the same workstation*.

NORMALIZED HYPERVOLUME REGRET

**63%**

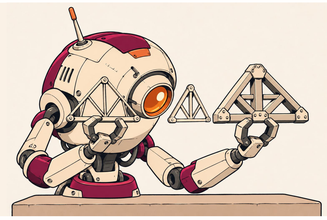
**lower than MOBO**

*Across 23 tasks*, mean regret is **0.043** for OffPSL and **0.117** for runner-up MOBO.

DESIGN MOTIVATION

## Which truss fits the new priority?

On a Pareto front, improving one objective requires a sacrifice in another. An engineer may favor **low mass** for one job and **low deflection** for the next.



### A lighter truss

Reducing material saves weight, but the structure may bend more.

### A stiffer truss

Adding reinforcement limits deflection at the cost of extra mass.

Historical tests pair designs with outcomes. The engineer supplies the preference later. With a pool of candidates, someone must search it again *each time the priorities change*.

**OffPSL returns one design for the weights the engineer specifies.**

OFFLINE TRAINING AND INFERENCE

## How a preference becomes a design

The model learns a map from objective weights to design coordinates. Two models fitted to the archive supply its training signal.

**Fixed archive of designs and measured objectives**

**Frozen surrogate ensemble**

Score each objective with $\mu + \beta\sigma$. Ensemble disagreement raises its predicted cost.

**Frozen denoising autoencoder**

**Large or unstable reconstruction errors** flag designs far from the offline examples.

**Optimize only the preference-to-design map**

### What the map minimizes

Smooth Tchebycheff scalarization weights each predicted objective against an ideal point just beyond the observed minima. Denoising penalties keep proposed designs near regions represented in the archive.

AT QUERY TIME

**Preference weights → Learned map → Box projection**

**One deterministic forward pass** returns a design. A sweep of weights produces a set of trade-offs with a known query for each design.

EVALUATION PROTOCOL

## 23 tasks, one fixed candidate budget

**5 ZDT**
Two-objective synthetic tasks

**7 DTLZ**
Scalable synthetic tasks

**11 RE**
Engineering design tasks

Each method submits **256 candidates per task and seed**. All candidates are projected into the same legal design box. The oracle evaluates the fixed set *only after submission*.

OffPSL, PCD and PGD-MOO use ten seeds. Four further baselines were rerun with five seeds. Published surrogate results also enter the main ranking.

**Hypervolume (HV) ↑**

Objective-space volume dominated by the candidate set relative to a fixed reference point.

**Family-balanced rank ↓**

Rank methods within each task, average within each family, then weight the three families equally.

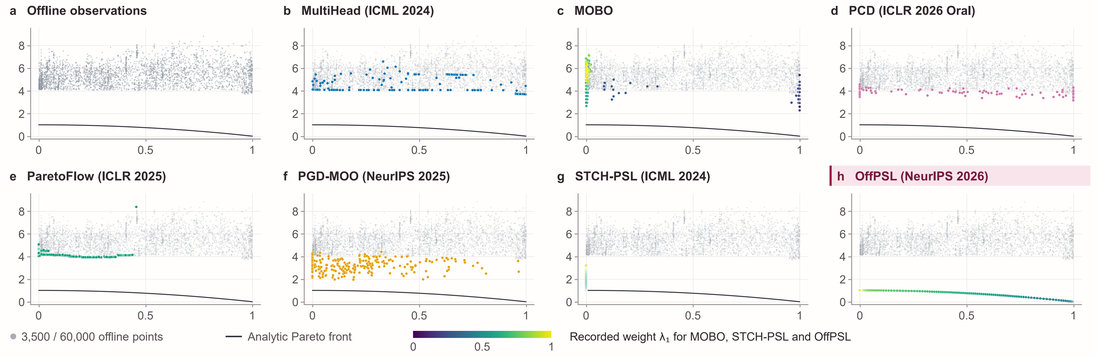
### How to read the 63% result

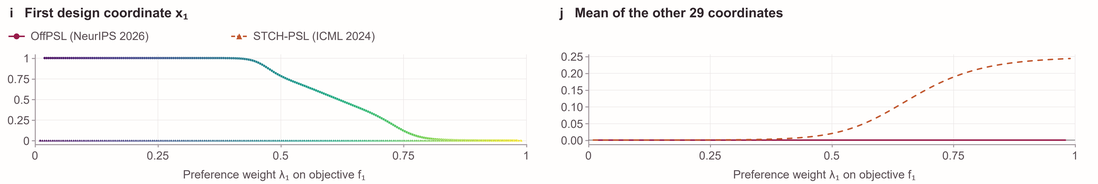
Regret scales from 0 for the best observed HV to 1 for the worst on each task. Across 23 tasks, OffPSL averages **0.043** versus **0.117** for MOBO.

ZDT2 COMPARISON / 30 DESIGN VARIABLES / SEED 1000

## OffPSL spans the Pareto front on ZDT2

Each method contributes **all 256 evaluated designs**. STCH-PSL stays at $f_1 < 0.001$, while OffPSL **follows the front from $f_1 \approx 0$ to 1**. The panels share *the same axes* ($f_1$ horizontal, $f_2$ vertical). Lower is better for both.




As $\lambda_1$ increases, OffPSL moves $x_1$ **from nearly one to nearly zero**. The mean of the other 29 coordinates stays **below $10^{-6}$** throughout the sweep, close to the Pareto-set value of zero.

RE22 / ARCHIVE COVERAGE

## Which trade-offs are missing?

Each training subset retains half of the archive. We change which trade-offs it contains and average the distance to the reference front *over five seeds*.

| Retained region | Mean IGD ↓ |
|---|---|
| **Uniform** | **0.004** |
| Extremes | 0.018 |
| One gap | 0.056 |
| Central only | 0.108 |

Uniform sampling gives **0.004 IGD**. Keeping only the centre raises it to **0.108** *at the same sample size*.

SUPPORT PRIOR / 23 TASKS

## Denoising versus kNN

Four support variants share one configuration across **23 tasks and five seeds** each (460 runs).

**17 / 23**
**tasks with higher HV**

**18 / 23**
**tasks with lower IGD**

Compared with no support term, denoising reduces near-duplicates from **1.9% to 0.8%** and clipped designs from **29% to 23%**.

ZDT2 / THREE RECORDED QUERIES

## What does a weight return?

| Weight $\lambda_1$ | Design $x_1$ | Objective $f_1$ | Objective $f_2$ |
|---|---|---|---|
| 0.251 | 1.000 | 1.000 | 0.000 |
| 0.498 | 0.790 | 0.790 | 0.375 |
| 0.749 | 0.123 | 0.123 | 0.985 |

OffPSL on ZDT2, seed 1000. Values are rounded to three decimals.

ACROSS 23 TASKS

## Best average HV rank in all three families

HV ranks compare **16 methods**. The last column ranks **seven methods** by front distance (IGD). Lower ranks are better for both metrics.

| Method (source) | ZDT | DTLZ | RE | HV avg. | W/T/L | IGD rank |
|---|---|---|---|---|---|---|
| **OffPSL (NeurIPS 2026)** | **4.30** | **3.29** | **3.86** | **3.82** | **—** | **3.13** |
| MultipleModels-IOM* (NeurIPS 2022) | 5.90 | 8.50 | 5.18 | 6.53 | 12/6/5 | — |
| MultipleModels-ICT* (NeurIPS 2023) | 4.80 | 10.50 | 4.91 | 6.74 | 12/5/6 | — |
| MultipleModels* (ICML 2024) | 5.60 | 9.64 | 5.05 | 6.76 | 15/2/6 | — |
| MultiHead-PCGrad* (NeurIPS 2020) | 7.00 | 7.86 | 7.00 | 7.29 | 16/3/4 | — |
| MOBO | 10.70 | 5.50 | 6.41 | 7.54 | 15/6/2 | 3.26 |
| MultiHead* (ICML 2024) | 5.80 | 11.07 | 7.14 | 8.00 | 15/3/5 | — |
| STCH-PSL (ICML 2024) | 7.20 | 7.07 | 10.05 | 8.11 | 17/3/3 | 4.30 |
| MultipleModels-RoMA* (NeurIPS 2021) | 7.70 | 8.36 | 8.95 | 8.34 | 16/3/4 | — |
| PGD-MOO (NeurIPS 2025) | 6.50 | 6.50 | 12.41 | 8.47 | 16/5/2 | 4.04 |
| PCD (ICLR 2026 Oral) | 10.50 | 5.29 | 10.68 | 8.82 | 17/5/1 | 3.30 |
| MultipleModels-TriMentoring* (NeurIPS 2023) | 10.60 | 10.57 | 5.64 | 8.94 | 16/4/3 | — |
| MultipleModels-COM* (ICML 2021) | 10.00 | 10.36 | 8.91 | 9.76 | 18/4/1 | — |
| ParetoFlow (ICLR 2025) | 12.60 | 8.43 | 13.14 | 11.39 | 22/0/1 | 4.13 |
| MultiHead-GradNorm* (ICML 2018) | 11.00 | 12.21 | 13.45 | 12.22 | 21/1/1 | — |
| PSL-MOBO (NeurIPS 2022) | 15.80 | 10.86 | 13.23 | 13.29 | 21/1/1 | 5.83 |

W/T/L counts OffPSL wins, ties and losses (*HV tie tolerance 0.01*). Asterisks mark published benchmark results.

Paired HV tests favor OffPSL against **all six rerun baselines** (*Holm-adjusted $p \le 0.0065$*).

### How to read the method labels

MultiHead and MultipleModels come from Off-MOO-Bench (ICML 2024). Each named variant cites its training method.
STCH-PSL and PSL-MOBO use offline adaptations.

COMPUTATION COST

## Training and generation time

Wall-clock seconds *on the same RTX 4070 Ti SUPER workstation*.

| Method (source) | Training (s) | Generate 256 (s) |
|---|---|---|
| **OffPSL (NeurIPS 2026)** | **262.97** | **0.06** |
| STCH-PSL (ICML 2024) | 12.37 | 0.04 |
| PSL-MOBO (NeurIPS 2022) | 15.20 | 0.03 |
| MultiHead-PCGrad (NeurIPS 2020) | 239.38 | 7.61 |
| MultipleModels-ICT (NeurIPS 2023) | 748.26 | 7.95 |
| PGD-MOO (NeurIPS 2025) | 2,961.81 | 7.83 |
| ParetoFlow (ICLR 2025) | 4,404.78 | 32.62 |
| MOBO | 0.78 | 3,309.40 |

After training, PGD-MOO takes **7.83 s** per 256 designs and OffPSL takes **0.06 s**, a roughly **130×** difference.

NATIVE PREFERENCE SWEEPS / SEVEN TWO-OBJECTIVE TASKS

## Smooth sweeps can still repeat designs

| Method (source) | Order viol. ↓ | Duplicates ↓ | Bins ↑ | Midpoint ↑ |
|---|---|---|---|---|
| **OffPSL (NeurIPS 2026)** | **0.141** | **7.2%** | **30.6** | **0.974** |
| STCH-PSL (ICML 2024) | 0.124 | 16.4% | 26.9 | 0.976 |
| PSL-MOBO (NeurIPS 2022) | 0.157 | 33.2% | 8.1 | 1.000 |
| MOBO | 0.445 | 0.1% | 26.6 | 0.218 |

OffPSL covers **30.6 bins** with **7.2% near-duplicates**. PSL-MOBO has a midpoint score of 1.000 but reaches only **8.1 bins**.

Order violations count reversals between adjacent weights. Bins measure breadth in objective space. Midpoints measure local continuity.