

# Novelty Assessment Report

**Paper:** Pareto-Conditioned Diffusion Models for Offline Multi-Objective Optimization

**PDF URL:** <https://openreview.net/pdf?id=S2Q00li155>

**Venue:** ICLR 2026 Conference Submission

**Year:** 2026

**Report Generated:** 2026-01-01

## Abstract

Multi-objective optimization (MOO) arises in many real-world applications where trade-offs between competing objectives must be carefully balanced. In the offline setting, where only a static dataset is available, the main challenge is generalizing beyond observed data. We introduce Pareto-Conditioned Diffusion (PCD), a novel framework that formulates offline MOO as a conditional sampling problem. By conditioning directly on desired trade-offs, PCD avoids the need for explicit surrogate models. To effectively explore the Pareto front, PCD employs a reweighting strategy that focuses on high-performing samples and a reference-direction mechanism to guide sampling towards novel, promising regions beyond the training data. Experiments on standard offline MOO benchmarks show that PCD achieves highly competitive performance and, importantly, demonstrates greater consistency across diverse tasks than existing offline MOO approaches.

### Disclaimer

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: mingzhang23@m.fudan.edu.cn

## Core Task Landscape

This paper addresses: **Offline Multi-Objective Optimization from Static Datasets**

A total of **50 papers** were analyzed and organized into a taxonomy with **16 categories**.

### Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Surrogate Modeling and Generative Approaches**
- **Reinforcement Learning Formulations**
- **Evolutionary and Metaheuristic Algorithms**
- **Direct Optimization and Ranking-Based Methods**
- **Application-Specific Offline MOO**
- **Methodological Foundations and Benchmarking**

### Complete Taxonomy Tree

- Offline Multi-Objective Optimization from Static Datasets Survey Taxonomy
- Surrogate Modeling and Generative Approaches
  - Regression-Based Surrogate Models (5 papers)
  - [1] Offline model-based optimization by learning to rank (Xue. Ke, 2024) [View paper](#)
  - [2] Offline Model-Based Optimization: Comprehensive Review (Kim, 2025) [View paper](#)
  - [6] Offline data-driven evolutionary optimization using selective surrogate ensembles (Handing Wang, 2018) [View paper](#)
  - [12] Offline data-driven multiobjective optimization: Knowledge transfer between surrogates and generation of final solutions (Cuie Yang, 2019) [View paper](#)
  - [41] Multifidelity surrogates-assisted multi-objective particle swarm algorithm for offline data-driven optimization (Yingying Cui, 2024) [View paper](#)
  - Generative Modeling Approaches ★ (4 papers)
  - [0] Pareto-Conditioned Diffusion Models for Offline Multi-Objective Optimization (Anon et al., 2026) [View paper](#)
  - [3] Paretoflow: Guided flows in multi-objective optimization (Yuan Ye, 2024) [View paper](#)
  - [9] Offline data-driven multiobjective optimization evolutionary algorithm based on generative adversarial network (Yu Zhang, 2022) [View paper](#)
  - [27] Preference-Guided Diffusion for Multi-Objective Offline Optimization (Annadani, 2025) [View paper](#)
- Reinforcement Learning Formulations
  - Policy-Regularized and Constrained Offline MORL (3 papers)
  - [14] Policy-regularized offline multi-objective reinforcement learning (Lin Qian, 2024) [View paper](#)
  - [18] Offline constrained multi-objective reinforcement learning via pessimistic dual value iteration (Runzhe Wu, 2021) [View paper](#)
  - [21] An offline adaptation framework for constrained multi-objective reinforcement learning (Qian Lin, 2024) [View paper](#)
  - Diffusion and Sequence Modeling in Offline MORL (2 papers)
  - [31] Prior-Guided Diffusion Planning for Offline Reinforcement Learning (Donghyeon Ki, 2025) [View paper](#)
  - [44] Multi-Objective Decision Transformers for Offline Reinforcement Learning (Ghanem, 2023) [View paper](#)
  - Preference Generalization and Pareto Front Learning (3 papers)
  - [33] MODULI: Unlocking Preference Generalization via Diffusion Models for Offline Multi-Objective Reinforcement Learning (Yuan, 2024) [View paper](#)
  - [36] Benchmarking Offline Multi-Objective Reinforcement Learning in Critical Care (Aryaman Bansal, 2025) [View paper](#)
  - [38] Scaling Pareto-Efficient Decision Making Via Offline Multi-Objective RL (Zhu, 2023) [View paper](#)
  - Hybrid Online-Offline and Multi-Task RL (3 papers)

- [37] Uni-O4: Unifying Online and Offline Deep Reinforcement Learning with Multi-Step On-Policy Optimization (Lei Kun, 2023) [View paper](#)
- [40] On multi-objective policy optimization as a tool for reinforcement learning: Case studies in offline RL and finetuning (Abdolmaleki, 2021) [View paper](#)
- [45] Offline Multi-Task Multi-Objective Data-Driven Evolutionary Algorithm with Language Surrogate Model and Implicit Q-Learning (Xian-Rong Zhang, 2025) [View paper](#)
- Evolutionary and Metaheuristic Algorithms
  - Genetic Algorithms for Offline MOO (2 papers)
  - [4] Genetic-programming-based multi-objective optimization of strategies for home energy-management systems (Jernej Zupančič, 2020) [View paper](#)
  - [8] Genetic algorithm-based multi-objective optimization model for software bugs prediction (Bakre Oluseye Musinat, 2022) [View paper](#)
  - Dynamic Multi-Objective Optimization (3 papers)
  - [5] A dynamic multi-objective optimization method based on classification strategies (Fei Wu, 2023) [View paper](#)
  - [26] Dynamic multi-objective optimization: a survey of the state-of-the-art (C. Raquel, 2013) [View paper](#)
  - [43] Dynamic multi-objective optimization and decision-making using modified NSGA-II: A case study on hydro-thermal power scheduling (Kalyanmoy Deb, 2007) [View paper](#)
  - NSGA-II and Metaheuristic Variants (3 papers)
  - [10] Research on coordinated control strategy of distributed static synchronous series compensator based on multi-objective optimization immune algorithm (Yu Wang, 2024) [View paper](#)
  - [20] Multi-Objective Process Parameter Optimization of Ultrasonic Rolling Combining Machine Learning and Non-Dominated Sorting Genetic Algorithm-II (Junying Chen, 2024) [View paper](#)
  - [29] A Multi-Objective Approach for Optimal Sizing and Placement of Distributed Generators and Distribution Static Compensators in a Distribution Network Using the Black Widow Optimization Algorithm (Rameez Shaikh, 2024) [View paper](#)
- Direct Optimization and Ranking-Based Methods (2 papers)
  - [7] Offline multi-objective optimization (Xue. Ke, 2024) [View paper](#)
  - [35] Offline Multi-Objective Bandits: From Logged Data to Pareto-Optimal Policies (J Cheng, 2025) [View paper](#)
- Application-Specific Offline MOO
  - Engineering Design and Manufacturing (7 papers)
  - [16] Enhanced data-driven topology design with multilevel mesh and correlation-based mutation for stress-related multi-objective optimization (Jun Yang, 2025) [View paper](#)
  - [19] Batch Bayesian optimization via multi-objective acquisition ensemble for automated analog circuit design (Wenlong Lyu, 2018) [View paper](#)
  - [22] Static and dynamic analyses and multi-objective optimization of wafer thinning machine's design variables (Guijiu Xie, 2024) [View paper](#)
  - [28] Effect of texture parameters on the lubrication performance of static and dynamic pressure thrust bearings and multi-objective optimization (Xiaodong Yu, 2024) [View paper](#)
  - [32] Multi objective optimization of composite laminate repaired by patches in considering static strength and fatigue life (Xin Yang, 2024) [View paper](#)
  - [34] Development of an integrated intelligent BIM-based model for multi-objective optimization in engineering assembly processes (Zhang, 2025) [View paper](#)
  - [47] Data-Driven Multi-Objective Optimization of Power Consumption and Quality for Finishing Rolling Process (Xiaolu Chen, 2024) [View paper](#)
  - Recommender Systems and Information Retrieval (3 papers)
  - [15] Controllable Multi-Objective Re-ranking with Policy Hypernetworks (Sirui Chen, 2023) [View paper](#)
  - [17] Offline Multi-Objective Optimization (OMOO) in Search Page Layout Optimization Using Off-policy Evaluation (Pratik Lahiri, 2024) [View paper](#)
  - [23] Static video summarization with multi-objective constrained optimization (M. Dhanushree, 2024) [View paper](#)
  - Robotics and Path Planning (3 papers)
  - [11] Modeling and Solving Multi-Objective Path Planning Problem for Cooperative Cable-Suspended Load Transportation Considering the Time Variable Risk (Amir Arslan Haghray, 2025) [View paper](#)
  - [13] Preference Based Multi-Objective Optimization Techniques for Coverage Path Planning of a Mobile Robot (Monex Sharma, 2023) [View paper](#)
  - [42] A mixed capability-based and optimization methodology for human-robot task allocation and scheduling (Andrea Monguzzi, 2022) [View paper](#)
  - Network and System Optimization (3 papers)
  - [24] MORM: A Multi-objective Optimized Replication Management strategy for cloud storage cluster (Saiqin Long, 2014) [View paper](#)
  - [48] Load-Aware Multi-Objective Optimization of Controller and Datastore Placement in Distributed Sdns (Kang Xingyuan, 2025) [View paper](#)
  - [49] Toward Autonomous and Efficient Cybersecurity: A Multi-Objective AutoML-based Intrusion Detection System (Li Yang, 2025) [View paper](#)
  - Process and Industrial Optimization (2 papers)
  - [46] Sustainability-Driven Decision Making: A Multi-Objective Optimisation Approach to Optimise Closed-Loop Supply Chains of Electrical Drives (Ashma, 2025) [View paper](#)
  - [50] Dynamic multi-objective optimization of batch chromatographic separation processes (A. Holmqvist, 2015) [View paper](#)
- Methodological Foundations and Benchmarking (3 papers)
  - [25] Multi-objective Optimization: A Case Study (Hasan, 2023) [View paper](#)
  - [30] Handling constrained multi-objective optimization problems with heterogeneous evaluation times: proof-of-principle results (Julian Blank, 2022) [View paper](#)
  - [39] Sistem Pendukung Keputusan Penerimaan Siswa Baru dengan Metode Multi-Objective Optimization by Ratio Analysis (Eka Prasetya Adhy Sugara, 2025) [View paper](#)

## Narrative

Core task: offline multi-objective optimization from static datasets. This field addresses the challenge of discovering Pareto-optimal solutions when objective evaluations are expensive or unavailable, relying instead on pre-collected data. The taxonomy reveals several

complementary methodological branches. Surrogate Modeling and Generative Approaches build learned models—ranging from Gaussian processes to deep generative networks—that approximate objective functions or directly synthesize candidate designs. Reinforcement Learning Formulations recast the search as a sequential decision problem, enabling policy-based exploration guided by logged trajectories. Evolutionary and Metaheuristic Algorithms adapt population-based search to leverage surrogate predictions, while Direct Optimization and Ranking-Based Methods focus on gradient-driven or preference-informed strategies. Application-Specific Offline MOO tailors these techniques to domains such as circuit design, energy management, and robotics, and Methodological Foundations and Benchmarking establishes theoretical guarantees and standardized testbeds.

Recent work highlights a tension between model fidelity and sample efficiency. Surrogate ensembles and multifidelity schemes balance accuracy with computational cost, while generative models like Paretoflow[3] and Preference Guided Diffusion[27] learn to sample diverse Pareto solutions directly from data. Pareto Conditioned Diffusion[0] sits within this generative modeling cluster, emphasizing conditional synthesis that respects multi-objective trade-offs without requiring online evaluations. Compared to GAN Based Offline[9], which also employs deep generative architectures, Pareto Conditioned Diffusion[0] leverages diffusion processes for more stable training and finer control over the generated front. Meanwhile, ranking-based approaches like Learning to Rank[1] offer an alternative by ordering candidates without explicit surrogate construction. These contrasting strategies reflect ongoing debates about whether to invest in high-fidelity surrogates, exploit flexible generative priors, or sidestep function approximation altogether through preference learning.

---

## Related Works in Same Category

The following **3 sibling papers** share the same taxonomy leaf node with the original paper:

### 1. Paretoflow: Guided flows in multi-objective optimization

**Authors:** Yuan Ye, Chen Can, Ye Yuan, Pal, Christopher, et al. (9 authors total) | **Year/Venue:** 2024 | **URL:** [View paper](#)

#### Abstract

In offline multi-objective optimization (MOO), we leverage an offline dataset of designs and their associated labels to simultaneously minimize multiple objectives. This setting more closely mirrors complex real-world problems compared to single-objective optimization. Recent works mainly employ evolutionary algorithms and Bayesian optimization, with limited attention given to the generative modeling capabilities inherent in such data. In this study, we explore generative modeling in offline MOO...

#### Relationship Analysis

Both papers belong to the Generative Modeling Approaches category, employing generative models to synthesize candidate solutions in offline multi-objective optimization. They share the core approach of using generative models (diffusion for PCD, flow matching for ParetoFlow) conditioned on objectives to generate Pareto-optimal solutions without explicit surrogate optimization loops. The key difference is that PCD uses diffusion models with a reweighting strategy and reference-direction mechanism for conditioning, while ParetoFlow uses flow matching with multi-objective predictor guidance and neighboring evolution to explore the Pareto front.

---

### 2. Offline data-driven multiobjective optimization evolutionary algorithm based on generative adversarial network

**Authors:** Yu Zhang, Wang Hu, Wen Yao, Li-Xian Lian, Gary G. Yen, et al. (7 authors total) | **Year/Venue:** 2022 | **URL:** [View paper](#)

#### Abstract

Usually, data-driven multiobjective optimization problems (DD-MOPs) are indirectly solved by evolutionary algorithms through the built surrogate model which is well-trained from sample data. However, in most DD-MOPs, only a few available data can be practically collected from real engineering experiments due to the unaffordable cost and time. The key challenge in such a DD-MOP is to prevent the serious deterioration on the accuracy of the obtained approximate Pareto front. In this article, two n...

#### Relationship Analysis

Both papers belong to the Generative Modeling Approaches category, employing generative models to address offline multi-objective optimization from static datasets. The original paper (PCD) uses diffusion models with Pareto-conditioned sampling and reweighting strategies to directly generate solutions without explicit surrogate models, while the candidate paper uses GANs for dual purposes: augmenting limited training data and providing critical fitness scores to guide evolutionary algorithms with surrogate models. The key difference is that PCD eliminates the need for separate surrogate predictors and evolutionary search, whereas the candidate paper integrates GANs into a traditional surrogate-based evolutionary framework.

---

### 3. Preference-Guided Diffusion for Multi-Objective Offline Optimization

**Authors:** Annadani, Yashas, Belakaria, Syrine, Yashas Annadani, et al. (15 authors total) | **Year/Venue:** 2025 • arXiv.org | **URL:** [View paper](#)

#### Abstract

Offline multi-objective optimization aims to identify Pareto-optimal solutions given a dataset of designs and their objective values. In this work, we propose a preference-guided diffusion model that generates Pareto-optimal designs by leveraging a classifier-based guidance mechanism. Our guidance classifier is a preference model trained to predict the probability that one design dominates another, directing the diffusion model toward optimal regions of the design space. Crucially, this preferen...

#### Relationship Analysis

Both papers belong to the Generative Modeling Approaches category, employing diffusion models for offline multi-objective optimization by generating candidate solutions conditioned on objectives. They overlap in using classifier-free guidance mechanisms to steer diffusion models toward high-quality regions and both address the challenge of generalizing beyond observed data to discover Pareto-optimal solutions. The key difference is that the candidate paper uses a preference-based classifier trained on pairwise dominance relationships to guide generation with explicit diversity-aware criteria, while the original paper (PCD) conditions directly on target objectives using a reweighting strategy and reference-direction mechanism without an explicit preference classifier.

---

## Contributions Analysis

**Overall novelty summary.** The paper introduces Pareto-Conditioned Diffusion (PCD), a generative framework that formulates offline multi-objective optimization as conditional sampling, avoiding explicit surrogate models. It resides in the 'Generative Modeling Approaches' leaf under 'Surrogate Modeling and Generative Approaches', alongside three sibling papers. This leaf represents a relatively sparse research direction within the broader taxonomy of fifty papers across approximately thirty-six topics, suggesting that generative modeling for offline MOO remains an emerging area compared to more established surrogate regression or evolutionary methods.

The taxonomy tree positions PCD within a branch that contrasts with 'Regression-Based Surrogate Models', which use ensembles or neural networks to approximate objectives, and 'Direct Optimization and Ranking-Based Methods', which bypass learned models entirely. Neighboring branches include 'Reinforcement Learning Formulations', which recast MOO as sequential decision-making, and 'Evolutionary and Metaheuristic Algorithms', which adapt population-based search. The scope note for PCD's leaf explicitly excludes regression surrogates, clarifying that generative approaches synthesize candidates rather than merely predicting objective values, distinguishing PCD from methods that rely on function approximation.

Among thirty candidates examined through limited semantic search, none clearly refute any of PCD's three contributions: the core framework, the multi-objective reweighting strategy, or the reference-direction mechanism. Each contribution was assessed against ten candidates, with zero refutable overlaps identified. This suggests that within the examined scope, PCD's combination of Pareto conditioning, reweighting for high-performing samples, and reference-direction guidance appears distinct. However, the analysis is constrained by the search scale and does not claim exhaustive coverage of all prior generative MOO work.

Given the limited search scope of thirty top-K semantic matches, the analysis indicates that PCD occupies a relatively novel position within generative offline MOO. The absence of refutable prior work among examined candidates, combined with the sparse population of its taxonomy leaf, suggests the approach introduces fresh mechanisms. Nonetheless, the findings reflect only the examined literature subset and do not preclude the existence of related work beyond the search boundary.

---

This paper presents **3 main contributions**, each analyzed against relevant prior work:

### **Contribution 1: Pareto-Conditioned Diffusion (PCD) framework**

**Description:** PCD reframes offline multi-objective optimization as a conditional sampling problem, enabling direct generation of high-quality solutions conditioned on target trade-offs without requiring explicit surrogate models or separate optimization algorithms. This provides a unified end-to-end approach that simplifies the optimization process.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

#### **1. Preference-Guided Diffusion for Multi-Objective Offline Optimization**

URL: [View paper](#)

##### **Brief Assessment**

Preference Guided Diffusion[27] uses classifier-based guidance with a preference model to direct diffusion, while PCD conditions directly on target objectives without explicit surrogate models. These are distinct technical approaches to offline multi-objective optimization.

---

#### **2. PILOT: equivariant diffusion for pocket-conditioned de novo ligand generation with multi-objective guidance via importance sampling**

URL: [View paper](#)

##### **Brief Assessment**

PILOT[53] focuses on structure-based drug design with pocket-conditioned ligand generation using multi-objective guidance via importance sampling, not on general offline multi-objective optimization as a conditional sampling problem.

---

#### **3. Shipgen: A diffusion model for parametric ship hull generation with multiple objectives and constraints**

URL: [View paper](#)

##### **Brief Assessment**

Shipgen[55] focuses on parametric ship hull generation with classifier guidance for specific engineering objectives (drag, volume), not on general offline multi-objective optimization through Pareto-conditioned sampling.

---

#### **4. Graph diffusion policy optimization**

URL: [View paper](#)

##### **Brief Assessment**

Graph Diffusion Policy[59] focuses on optimizing graph diffusion models for graph generation tasks using reinforcement learning, not on multi-objective optimization with Pareto conditioning for general offline optimization problems.

---

#### **5. Addressing high-performance data sparsity in metasurface inverse design using multi-objective optimization and diffusion probabilistic models.**

URL: [View paper](#)

##### **Brief Assessment**

Metasurface Inverse Design[54] focuses on metasurface design using diffusion models with attention mechanisms, not on offline multi-objective optimization as a conditional sampling problem for general design tasks.

---

#### **6. DyMO: Training-Free Diffusion Model Alignment with Dynamic Multi-Objective Scheduling**

URL: [View paper](#)

##### **Brief Assessment**

DyMO[57] focuses on text-to-image diffusion model alignment with human preferences through training-free guidance during inference, not on offline multi-objective optimization problems. The candidate addresses a fundamentally different problem domain (image generation alignment) compared to the original's focus on multi-objective optimization with Pareto fronts.

---

#### **7. Airfoil-DDPM: A flexible airfoil generative design method using a multi-objective sampling based diffusion model**

URL: [View paper](#)

##### **Brief Assessment**

Airfoil-DDPM[56] focuses on airfoil design using conditional denoising diffusion models for multi-objective problems, but does not demonstrate the specific PCD framework elements (reweighting strategy, reference-direction mechanism, or direct conditioning on target trade-offs without surrogate models).

---

#### **8. EmoDM: A Diffusion Model for Evolutionary Multi-objective Optimization**

URL: [View paper](#)

##### **Brief Assessment**

EmoDM[51] focuses on learning evolutionary search patterns from previously solved tasks for generalization to new problems, rather than conditioning on target trade-offs for offline optimization as in PCD.

---

#### **9. A reward-directed diffusion framework for generative design optimization**

URL: [View paper](#)

##### **Brief Assessment**

Reward Directed Diffusion[52] focuses on single-objective optimization using reward-directed sampling for engineering design (ship hulls, airfoils), not multi-objective optimization with Pareto conditioning. The candidate uses reward functions to guide diffusion toward

high-performance designs in a single-objective context, whereas PCD addresses the fundamentally different problem of balancing multiple conflicting objectives along a Pareto front.

---

## 10. Protein design with guided discrete diffusion

URL: [View paper](#)

### Brief Assessment

Protein Discrete Diffusion[58] focuses on protein sequence design using discrete diffusion with gradient guidance in hidden states, not on multi-objective optimization with Pareto conditioning. The candidate addresses single or multi-objective protein design through guided sampling, while the original paper specifically reframes offline MOO as conditional sampling on target trade-offs without explicit surrogate models.

---

## Contribution 2: Multi-objective reweighting strategy

**Description:** A reweighting strategy based on dominance numbers that emphasizes high-performing samples near the Pareto front during training. This allows the model to generalize more accurately in regions containing well-performing solutions while reducing emphasis on low-performing areas.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

## 1. An integrated TOPSIS and ARAS method multi-criteria decision-making approach for optimizing investment portfolios using goal programming and genetic algorithm

URL: [View paper](#)

### Brief Assessment

TOPSIS ARAS Portfolio[64] focuses on portfolio optimization using multi-criteria decision-making methods (TOPSIS, ARAS) combined with goal programming and genetic algorithms. It does not address reweighting strategies for Pareto front regions in multi-objective optimization, which is the core novelty of the original paper's contribution.

---

## 2. Multi-objective optimization

URL: [View paper](#)

### Brief Assessment

Multi-Objective Optimization[61] focuses on general multi-objective optimization theory and evolutionary algorithms, not on reweighting strategies for emphasizing Pareto front regions in offline settings or diffusion-based generative models.

---

## 3. Application and analysis of methods for selecting an optimal solution from the Pareto-optimal front obtained by multiobjective optimization

URL: [View paper](#)

### Brief Assessment

Pareto Solution Selection[60] focuses on selecting a single optimal solution from an already-computed Pareto front using decision-making methods (TOPSIS, GRA, SAW). It does not address training strategies or reweighting samples during model learning, which is the core of the original paper's contribution.

---

## 4. What weights work for you? Adapting weights for any Pareto front shape in decomposition-based evolutionary multiobjective optimisation

URL: [View paper](#)

### Brief Assessment

Adapting Weights[62] focuses on adapting weight vector distributions in decomposition-based evolutionary algorithms for multi-objective optimization, not on reweighting training samples based on dominance numbers to emphasize high-performing regions near the Pareto front during neural network training.

---

## 5. Multi-objective optimization for dynamic logistics scheduling based on hierarchical deep reinforcement learning

URL: [View paper](#)

### Brief Assessment

Logistics Hierarchical DRL[69] focuses on hierarchical reinforcement learning for logistics scheduling with adaptive weighting of multiple objectives (time, cost, quality), not on reweighting training samples based on dominance numbers to emphasize Pareto front regions during model training as in the original paper.

---

## 6. EIT Reconstruction Based on Pareto Multi-Objective Optimization

URL: [View paper](#)

### Brief Assessment

EIT Reconstruction[65] mentions reweighting in the context of electrical impedance tomography reconstruction, not offline multi-objective optimization. The candidate focuses on image reconstruction using Pareto optimization, which is a different application domain than the original paper's offline MOO framework.

---

## 7. Analysis of weighting and selection methods for pareto-optimal solutions of multiobjective optimization in chemical engineering applications

URL: [View paper](#)

### Brief Assessment

The candidate paper (Weighting Selection Analysis[63]) focuses on analyzing weighting and selection methods for Pareto-optimal solutions in chemical engineering applications, not on developing reweighting strategies for training generative models in offline multi-objective optimization.

---

## 8. Learning to optimize multi-objective alignment through dynamic reward weighting

URL: [View paper](#)

### Brief Assessment

Dynamic Reward Weighting[67] focuses on online RL for LLM alignment with dynamic weight adaptation during training, while the original paper addresses offline multi-objective optimization using static reweighting based on dominance numbers to emphasize Pareto front regions in a pre-collected dataset.

---

## 9. Searching for the Pareto frontier in multi-objective protein design

URL: [View paper](#)

### Brief Assessment

Pareto Protein Design[66] focuses on protein design with Pareto optimization but does not describe a reweighting strategy based on dominance numbers for emphasizing high-performing samples near the Pareto front during training.

---

## 10. Multi-Objective Optimization and Multi-Criteria Decision-Making Approach to Design a Multi-Tubular Packed-Bed Membrane Reactor in Oxidative Dehydrogenation of Ethane

URL: [View paper](#)

### Brief Assessment

Membrane Reactor Design[68] focuses on chemical engineering optimization using NSGA-II and MCDM methods for reactor design, not on reweighting strategies for machine learning models near Pareto fronts.

---

## Contribution 3: Reference-direction mechanism for conditioning

**Description:** A two-stage procedure for generating diverse and high-quality conditioning points that guide sampling toward novel, promising regions. The mechanism partitions the objective space using direction vectors and extrapolates representative points to enable exploration beyond the training data.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

### 1. From source to target and back: symmetric bi-directional adaptive gan

URL: [View paper](#)

#### Brief Assessment

Symmetric Bidirectional GAN[71] focuses on image-to-image domain adaptation using symmetric bi-directional generators for visual domain transfer, not on multi-objective optimization with reference direction mechanisms for generating conditioning points beyond observed data.

---

### 2. Dual-channel target speaker extraction based on conditional variational autoencoder and directional information

URL: [View paper](#)

#### Brief Assessment

Target Speaker Extraction[73] focuses on audio signal processing using directional information (DOA) for speaker extraction, not on generating diverse conditioning points for multi-objective optimization in offline settings.

---

### 3. MetaIP: Meta-Network-Based Intra Prediction With Customized Parameters for Video Coding

URL: [View paper](#)

#### Brief Assessment

MetaIP[72] focuses on video coding intra prediction with meta-networks that generate customized parameters for image blocks, not on multi-objective optimization or reference-direction mechanisms for generating diverse conditioning points beyond observed data.

---

### 4. Directing Nanoparticle Organization in Response to Diverse Chemical Inputs.

URL: [View paper](#)

#### Brief Assessment

Nanoparticle Chemical Inputs[76] focuses on biological signaling cascades and nanoparticle organization in DNA frameworks, not on reference direction mechanisms for generating conditioning points in optimization problems.

---

### 5. Heading direction with respect to a reference point modulates place-cell activity

URL: [View paper](#)

#### Brief Assessment

Heading Direction Modulation[74] studies neural activity modulation in hippocampal place cells based on heading direction relative to reference points in spatial navigation tasks. This is fundamentally different from the original paper's reference-direction mechanism for generating diverse conditioning points in offline multi-objective optimization through direction vectors and extrapolation.

---

### 6. Whole-body central processing of lateral line inputs encodes flow direction relative to the center-of-mass

URL: [View paper](#)

#### Brief Assessment

Lateral Line Processing[77] focuses on neural encoding of fluid flow direction in zebrafish lateral line sensory systems, not on generating diverse conditioning points for multi-objective optimization.

---

### 7. Rdfinet: reference-guided directional diverse face inpainting network

URL: [View paper](#)

#### Brief Assessment

Rdfinet[75] focuses on face image inpainting using reference face images to guide style transfer, not on multi-objective optimization or generating diverse conditioning points for exploring Pareto fronts in offline optimization settings.

---

### 8. Arithmetic Control of LLMs for Diverse User Preferences: Directional Preference Alignment with Multi-Objective Rewards

URL: [View paper](#)

#### Brief Assessment

Directional Preference Alignment[79] focuses on multi-objective reward modeling for LLM alignment with user preferences, not on offline multi-objective optimization with diffusion models for design problems.

---

### 9. Attack Deterministic Conditional Image Generative Models for Diverse and Controllable Generation

URL: [View paper](#)

#### Brief Assessment

Attack Conditional GAN[70] focuses on adversarial perturbations to deterministic image generation models for diversity, not on reference-direction mechanisms for multi-objective optimization or conditioning point generation in offline settings.

---

## 10. Redundancy parameterization and inverse kinematics of 7-DOF revolute manipulators

URL: [View paper](#)

### Brief Assessment

Redundancy Inverse Kinematics[78] addresses inverse kinematics for 7-DOF manipulators using reference direction functions for redundancy parameterization, not multi-objective optimization or generating diverse conditioning points for diffusion models.

---

## Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

---

## References

- [0] Pareto-Conditioned Diffusion Models for Offline Multi-Objective Optimization [View paper](#)
- [1] Offline model-based optimization by learning to rank [View paper](#)
- [2] Offline Model-Based Optimization: Comprehensive Review [View paper](#)
- [3] Paretoflow: Guided flows in multi-objective optimization [View paper](#)
- [4] Genetic-programming-based multi-objective optimization of strategies for home energy-management systems [View paper](#)
- [5] A dynamic multi-objective optimization method based on classification strategies [View paper](#)
- [6] Offline data-driven evolutionary optimization using selective surrogate ensembles [View paper](#)
- [7] Offline multi-objective optimization [View paper](#)
- [8] Genetic algorithm-based multi-objective optimization model for software bugs prediction [View paper](#)
- [9] Offline data-driven multiobjective optimization evolutionary algorithm based on generative adversarial network [View paper](#)
- [10] Research on coordinated control strategy of distributed static synchronous series compensator based on multi-objective optimization immune algorithm [View paper](#)
- [11] Modeling and Solving Multi-Objective Path Planning Problem for Cooperative Cable-Suspended Load Transportation Considering the Time Variable Risk [View paper](#)
- [12] Offline data-driven multiobjective optimization: Knowledge transfer between surrogates and generation of final solutions [View paper](#)
- [13] Preference Based Multi-Objective Optimization Techniques for Coverage Path Planning of a Mobile Robot [View paper](#)
- [14] Policy-regularized offline multi-objective reinforcement learning [View paper](#)
- [15] Controllable Multi-Objective Re-ranking with Policy Hypernetworks [View paper](#)
- [16] Enhanced data-driven topology design with multilevel mesh and correlation-based mutation for stress-related multi-objective optimization [View paper](#)
- [17] Offline Multi-Objective Optimization (OMOO) in Search Page Layout Optimization Using Off-policy Evaluation [View paper](#)
- [18] Offline constrained multi-objective reinforcement learning via pessimistic dual value iteration [View paper](#)
- [19] Batch Bayesian optimization via multi-objective acquisition ensemble for automated analog circuit design [View paper](#)
- [20] Multi-Objective Process Parameter Optimization of Ultrasonic Rolling Combining Machine Learning and Non-Dominated Sorting Genetic Algorithm-II [View paper](#)
- [21] An offline adaptation framework for constrained multi-objective reinforcement learning [View paper](#)
- [22] Static and dynamic analyses and multi-objective optimization of wafer thinning machine's design variables [View paper](#)
- [23] Static video summarization with multi-objective constrained optimization [View paper](#)
- [24] MORM: A Multi-objective Optimized Replication Management strategy for cloud storage cluster [View paper](#)
- [25] Multi-objective Optimization: A Case Study [View paper](#)
- [26] Dynamic multi-objective optimization: a survey of the state-of-the-art [View paper](#)
- [27] Preference-Guided Diffusion for Multi-Objective Offline Optimization [View paper](#)
- [28] Effect of texture parameters on the lubrication performance of static and dynamic pressure thrust bearings and multi-objective optimization [View paper](#)
- [29] A Multi-Objective Approach for Optimal Sizing and Placement of Distributed Generators and Distribution Static Compensators in a Distribution Network Using the Black Widow Optimization Algorithm [View paper](#)
- [30] Handling constrained multi-objective optimization problems with heterogeneous evaluation times: proof-of-principle results [View paper](#)
- [31] Prior-Guided Diffusion Planning for Offline Reinforcement Learning [View paper](#)
- [32] Multi objective optimization of composite laminate repaired by patches in considering static strength and fatigue life [View paper](#)
- [33] MODULI: Unlocking Preference Generalization via Diffusion Models for Offline Multi-Objective Reinforcement Learning [View paper](#)
- [34] Development of an integrated intelligent BIM-based model for multi-objective optimization in engineering assembly processes [View paper](#)
- [35] Offline Multi-Objective Bandits: From Logged Data to Pareto-Optimal Policies [View paper](#)
- [36] Benchmarking Offline Multi-Objective Reinforcement Learning in Critical Care [View paper](#)
- [37] Uni-O4: Unifying Online and Offline Deep Reinforcement Learning with Multi-Step On-Policy Optimization [View paper](#)
- [38] Scaling Pareto-Efficient Decision Making Via Offline Multi-Objective RL [View paper](#)
- [39] Sistem Pendukung Keputusan Penerimaan Siswa Baru dengan Metode Multi-Objective Optimization by Ratio Analysis [View paper](#)
- [40] On multi-objective policy optimization as a tool for reinforcement learning: Case studies in offline RL and finetuning [View paper](#)
- [41] Multifidelity surrogates-assisted multi-objective particle swarm algorithm for offline data-driven optimization [View paper](#)
- [42] A mixed capability-based and optimization methodology for human-robot task allocation and scheduling [View paper](#)
- [43] Dynamic multi-objective optimization and decision-making using modified NSGA-II: A case study on hydro-thermal power scheduling [View paper](#)
- [44] Multi-Objective Decision Transformers for Offline Reinforcement Learning [View paper](#)
- [45] Offline Multi-Task Multi-Objective Data-Driven Evolutionary Algorithm with Language Surrogate Model and Implicit Q-Learning [View paper](#)
- [46] Sustainability-Driven Decision Making: A Multi-Objective Optimisation Approach to Optimise Closed-Loop Supply Chains of Electrical Drives [View paper](#)
- [47] Data-Driven Multi-Objective Optimization of Power Consumption and Quality for Finishing Rolling Process [View paper](#)

- [48] Load-Aware Multi-Objective Optimization of Controller and Datastore Placement in Distributed Sdns [View paper](#)
- [49] Toward Autonomous and Efficient Cybersecurity: A Multi-Objective AutoML-based Intrusion Detection System [View paper](#)
- [50] Dynamic multi-objective optimization of batch chromatographic separation processes [View paper](#)
- [51] EmoDM: A Diffusion Model for Evolutionary Multi-objective Optimization [View paper](#)
- [52] A reward-directed diffusion framework for generative design optimization [View paper](#)
- [53] PILOT: equivariant diffusion for pocket-conditioned de novo ligand generation with multi-objective guidance via importance sampling [View paper](#)
- [54] Addressing high-performance data sparsity in metasurface inverse design using multi-objective optimization and diffusion probabilistic models. [View paper](#)
- [55] Shipgen: A diffusion model for parametric ship hull generation with multiple objectives and constraints [View paper](#)
- [56] Airfoil-DDPM: A flexible airfoil generative design method using a multi-objective sampling based diffusion model [View paper](#)
- [57] DyMO: Training-Free Diffusion Model Alignment with Dynamic Multi-Objective Scheduling [View paper](#)
- [58] Protein design with guided discrete diffusion [View paper](#)
- [59] Graph diffusion policy optimization [View paper](#)
- [60] Application and analysis of methods for selecting an optimal solution from the Pareto-optimal front obtained by multiobjective optimization [View paper](#)
- [61] Multi-objective optimization [View paper](#)
- [62] What weights work for you? Adapting weights for any Pareto front shape in decomposition-based evolutionary multiobjective optimisation [View paper](#)
- [63] Analysis of weighting and selection methods for pareto-optimal solutions of multiobjective optimization in chemical engineering applications [View paper](#)
- [64] An integrated TOPSIS and ARAS method multi-criteria decision-making approach for optimizing investment portfolios using goal programming and genetic algorithm [View paper](#)
- [65] EIT Reconstruction Based on Pareto Multi-Objective Optimization [View paper](#)
- [66] Searching for the Pareto frontier in multi-objective protein design [View paper](#)
- [67] Learning to optimize multi-objective alignment through dynamic reward weighting [View paper](#)
- [68] Multi-Objective Optimization and Multi-Criteria Decision-Making Approach to Design a Multi-Tubular Packed-Bed Membrane Reactor in Oxidative Dehydrogenation of Ethane [View paper](#)
- [69] Multi-objective optimization for dynamic logistics scheduling based on hierarchical deep reinforcement learning [View paper](#)
- [70] Attack Deterministic Conditional Image Generative Models for Diverse and Controllable Generation [View paper](#)
- [71] From source to target and back: symmetric bi-directional adaptive gan [View paper](#)
- [72] MetaIP: Meta-Network-Based Intra Prediction With Customized Parameters for Video Coding [View paper](#)
- [73] Dual-channel target speaker extraction based on conditional variational autoencoder and directional information [View paper](#)
- [74] Heading direction with respect to a reference point modulates place-cell activity [View paper](#)
- [75] Rdffinet: reference-guided directional diverse face inpainting network [View paper](#)
- [76] Directing Nanoparticle Organization in Response to Diverse Chemical Inputs. [View paper](#)
- [77] Whole-body central processing of lateral line inputs encodes flow direction relative to the center-of-mass [View paper](#)
- [78] Redundancy parameterization and inverse kinematics of 7-DOF revolute manipulators [View paper](#)
- [79] Arithmetic Control of LLMs for Diverse User Preferences: Directional Preference Alignment with Multi-Objective Rewards [View paper](#)