Novelty Assessment Report

Paper: Universal Inverse Distillation for Matching Models with Real-Data Supervision (No GANs)
PDF URL: https://openreview.net/pdf?id=8NuN5UzXLC
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2026-01-08

Abstract

While achieving exceptional generative quality, modern diffusion, flow, and other matching models suffer from slow inference, as they require many steps of iterative generation. Recent distillation methods address this by training efficient one-step generators under the guidance of a pre-trained teacher model. However, these methods are often constrained to only one specific framework, e.g., only to diffusion or only to flow models. Furthermore, these methods are naturally data-free, and to benefit from the usage of real data, it is required to use an additional complex adversarial training with an extra discriminator model. In this paper, we present \textbf{RealUID}, a unified distillation framework for all matching models that seamlessly incorporates real data into the distillation procedure without GANs. Our \textbf{RealUID} approach offers a simple theoretical foundation that covers previous distillation methods for Flow Matching and Diffusion models, and is also extended to their modifications, such as Bridge Matching and Stochastic Interpolants.

Core Task Landscape

This paper addresses: **Distilling Matching Models into Efficient One-Step Generators**

A total of **36 papers** were analyzed and organized into a taxonomy with **21 categories**.

Taxonomy Overview

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

- **Distribution Matching Distillation Methods**
- **Score Identity and Implicit Matching Distillation**
- **Adversarial and Hybrid Distillation Methods**
- **Consistency and Flow Trajectory Distillation**
- **Specialized Distillation Frameworks**
- **Unified and Theoretical Distillation Frameworks**
- **Domain-Specific Distillation Applications**

Complete Taxonomy Tree

- Distilling Matching Models into Efficient One-Step Generators Survey Taxonomy
- Distribution Matching Distillation Methods
  - KL Divergence-Based Distribution Matching (3 papers)
  - [4] One-Step Diffusion with Distribution Matching Distillation (Tianwei Yin, 2024) View paper
  - [24] Phased DMD: Few-step Distribution Matching Distillation via Score Matching within Subintervals (Fan Xiang-yu, 2025) View paper
  - [32] Improved Distribution Matching Distillation for Fast Image Synthesis (FrÃ©do Durand, 2024) View paper
  - Generalized Divergence Distribution Matching (3 papers)
  - [26] One-step Diffusion Models with F-Divergence Distribution Matching (Xu, 2025) View paper
  - [27] VarDiU: A Variational Diffusive Upper Bound for One-Step Diffusion Distillation (Wang Leyang, 2025) View paper
  - [31] One-step Diffusion Models with Bregman Density Ratio Matching (Zhu Yuanzhi, 2025) View paper
  - Unpaired and Bridge Matching Distillation (2 papers)
  - [12] Inverse Bridge Matching Distillation (Gushchin, 2025) View paper
  - [25] Regularized Distribution Matching Distillation for One-step Unpaired Image-to-Image Translation (Rakitin, 2024) View paper
- Score Identity and Implicit Matching Distillation
  - Score Identity Distillation (2 papers)
  - [2] Score identity Distillation: Exponentially Fast Distillation of Pretrained Diffusion Models for One-Step Generation (Zhou, 2024) View paper
  - [14] Few-step diffusion via score identity distillation (Zhou, 2025) View paper
  - Score Implicit Matching (1 papers)
  - [19] One-Step Diffusion Distillation through Score Implicit Matching (Zhengyang Geng, 2024) View paper
- Adversarial and Hybrid Distillation Methods
  - Adversarial Distribution Matching (2 papers)
  - [3] Adversarial distribution matching for diffusion distillation towards efficient image and video synthesis (LU Yanzuo, 2025) View paper
  - [20] Adversarial score identity distillation: Rapidly surpassing the teacher in one step (Zhou, 2024) View paper
  - GAN-Based Distillation (1 papers)
  - [17] Distilling Diffusion Models into Conditional GANs (Kang, 2024) View paper
  - Hybrid Flow-GAN and Multi-Objective Methods (2 papers)
  - [8] Sana-sprint: One-step diffusion with continuous-time consistency distillation (Chen Jun-song, 2025) View paper

- [23] Flow2GAN: Hybrid Flow Matching and GAN with Multi-Resolution Network for Few-step High-Fidelity Audio Generation (Zengwei Yao, 2025) View paper
- Consistency and Flow Trajectory Distillation
  - Consistency Distillation (2 papers)
  - [6] Self-Corrected Flow Distillation for Consistent One-Step and Few-Step Image Generation (Quan Dao, 2025) View paper
  - [34] Robust One-step Speech Enhancement via Consistency Distillation (Xu Liang, 2025) View paper
  - Trajectory Distribution Matching (1 papers)
  - [10] Learning Few-Step Diffusion Models by Trajectory Distribution Matching (Luo, 2025) View paper
  - Flow Generator Matching (1 papers)
  - [21] Flow Generator Matching (Huang Ze-min, 2024) View paper
- Specialized Distillation Frameworks
  - Multi-Student and Multi-Step Distillation (2 papers)
  - [5] Multi-student Diffusion Distillation for Better One-step Generators (Song YanKe, 2024) View paper
  - [16] Towards One-step Causal Video Generation via Adversarial Self-Distillation (Yang Yongqi, 2025) View paper
  - Variational and Expectation-Maximization Distillation (2 papers)
  - [11] SwiftBrush: One-Step Text-to-Image Diffusion Model with Variational Score Distillation (Thuan Hoang Nguyen, 2024) View paper
  - [33] EM Distillation for One-step Diffusion Models (Ruiqi Gao, 2024) View paper
  - Autoregressive Model Distillation (1 papers)
  - [9] Distilled Decoding 1: One-step Sampling of Image Auto-regressive Models with Flow Matching (Liu, 2024) View paper
- Unified and Theoretical Distillation Frameworks
  - Unified Distillation Frameworks ★ (2 papers)
  - [0] Universal Inverse Distillation for Matching Models with Real-Data Supervision (No GANs) (Anon et al., 2026) View paper
  - [36] On Distilling Generator Matching Models (Shankar, n.d.) View paper
  - Theoretical Foundations and Revisited Paradigms (2 papers)
  - [18] Diffusion Models Are Innate One-Step Generators (Zheng Bowen, 2024) View paper
  - [22] Revisiting Diffusion Models: From Generative Pre-training to One-Step Generation (Zheng Bowen, 2025) View paper
  - Simple and Regularized Distillation (2 papers)
  - [15] Simple Distillation for One-Step Diffusion Models (H Zhu, 2025) View paper
  - [35] One-step Flow Matching Generators (Z Huang, n.d.) View paper
- Domain-Specific Distillation Applications
  - Video and Temporal Generation Distillation (1 papers)
  - [7] VideoScene: Distilling Video Diffusion Model to Generate 3D Scenes in One Step (Wang Hanyang, 2025) View paper
  - Super-Resolution Distillation (1 papers)
  - [13] One diffusion step to real-world super-resolution via flow trajectory distillation (Li, 2025) View paper
  - Guided Text-to-Image Distillation (1 papers)
  - [29] Guided Score identity Distillation for Data-Free One-Step Text-to-Image Generation (Zhou, 2024) View paper
  - Latent Space and Representation Distillation (3 papers)
  - [1] Di o: Distilling masked diffusion models into one-step generator (Y Zhu, 2025) View paper
  - [28] MeanFlow Transformers with Representation Autoencoders (Zheyuan Hu, 2025) View paper
  - [30] ProtFlow: Fast Protein Sequence Design via Flow Matching on Compressed Protein Language Model Embeddings (Zhu Yiheng, 2025) View paper

Narrative

Core task: Distilling matching models into efficient one-step generators. The field addresses the challenge of compressing multi-step diffusion or flow models into fast, single-step samplers while preserving generation quality. The taxonomy reveals several complementary strategies: Distribution Matching Distillation Methods (e.g., Distribution Matching Distillation[4], Improved Distribution Matching[32]) focus on aligning student and teacher output distributions directly; Score Identity and Implicit Matching Distillation approaches (e.g., Score Identity Distillation[2], Score Implicit Matching[19]) leverage score-based objectives to guide the distillation; Adversarial and Hybrid Distillation Methods (e.g., Adversarial Distribution Matching[3], Adversarial Score Identity[20]) incorporate discriminative signals; Consistency and Flow Trajectory Distillation (e.g., Flow Trajectory Distillation[13], Self-Corrected Flow[6]) emphasize trajectory-level consistency; Specialized Distillation Frameworks target domain-specific constraints (e.g., VideoScene[7], Robust Speech Enhancement[34]); Unified and Theoretical Distillation Frameworks aim to provide general principles (e.g., Universal Inverse Distillation[0], Distilling Generator Matching[36]); and Domain-Specific Distillation Applications adapt these ideas to particular modalities.

Recent work explores trade-offs between distillation objectives, computational cost, and sample quality. Distribution matching methods like DMD[4] and its variants (Phased DMD[24], Regularized Distribution Matching[25]) offer simplicity but may struggle with mode coverage, while score-based approaches (Score Identity Distillation[2], Few-step Score Identity[14]) provide stronger theoretical grounding at the expense of additional computation. Adversarial techniques (Adversarial Distribution Matching[3], Flow2GAN[23]) can sharpen outputs but risk training instability. Universal Inverse Distillation[0] sits within the Unified and Theoretical Distillation Frameworks branch alongside Distilling Generator Matching[36], emphasizing a principled, general formulation that unifies multiple distillation paradigms. Compared to specialized methods like SwiftBrush[11] or domain-targeted frameworks, Universal Inverse Distillation[0] aims for broader applicability across generative tasks, bridging theoretical insights with practical one-step generation.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. On Distilling Generator Matching Models

Authors: S Shankar | **URL**: View paper

Abstract

â¦ Generator Matchingâ (IGM), a general framework for one-step distillation of generator matching modelsâ¦ Our method generalizes the recently proposed one-step diffusion distillation (â¦

Relationship Analysis

Both papers belong to the unified distillation frameworks category, aiming to distill matching models into efficient one-step generators. They overlap in addressing the slow inference problem of diffusion and flow matching models through distillation with real-data supervision. The key difference is that the original paper (RealUID) focuses on incorporating real data directly into the distillation loss without GANs through a linearization technique and inverse optimization framework, while the candidate paper (IGM) proposes a general implicit generator matching framework that uses an auxiliary model to approximate the implicit generator and extends to jump processes beyond standard diffusion/flow models.

Contributions Analysis

Overall novelty summary. The paper proposes RealUID, a unified distillation framework that compresses multi-step matching models (diffusion, flow, bridge matching, stochastic interpolants) into efficient one-step generators while incorporating real data without adversarial training. It resides in the 'Unified Distillation Frameworks' leaf alongside one sibling paper (Distilling Generator Matching). This leaf is part of the broader 'Unified and Theoretical Distillation Frameworks' branch, which contains only three leaves and seven papers total. The positioning suggests a relatively sparse research direction focused on general-purpose, theoretically grounded distillation methods rather than model-specific or domain-specific approaches.

The taxonomy reveals that most distillation research clusters around specialized strategies: Distribution Matching (three leaves, five papers), Score Identity methods (two leaves, three papers), Adversarial/Hybrid approaches (three leaves, five papers), and Consistency/Trajectory methods (three leaves, three papers). RealUID's unified framework contrasts with these narrower paradigms by claiming applicability across diffusion, flow, and bridge matching models. The 'Unified and Theoretical Distillation Frameworks' branch sits apart from domain-specific applications (video, super-resolution, text-to-image) and specialized frameworks (multi-student, variational, autoregressive), indicating the paper targets foundational methodology rather than task-specific optimization.

Among 30 candidates examined, the contribution-level analysis shows mixed novelty signals. The 'RealUID framework' and 'real data incorporation without GANs' contributions each examined 10 candidates with zero refutable overlaps, suggesting these aspects appear relatively novel within the limited search scope. However, the 'Universal Matching loss with real data (RealUM)' contribution found one refutable candidate among 10 examined, indicating at least one prior work in the candidate pool addresses overlapping ideas. The single sibling paper in the same taxonomy leaf likely represents closely related unified distillation work, though the analysis does not specify whether it was among the refutable candidates.

Based on the limited 30-candidate search, the work appears to occupy a less-crowded research direction (unified frameworks) compared to specialized distillation methods. The claim of unifying multiple matching model families without GANs shows some novelty signals, though one contribution has identifiable prior overlap. The analysis covers top-K semantic matches and does not constitute an exhaustive literature review, so additional related work may exist beyond the examined scope.

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

Contribution 1: RealUID: Universal distillation framework for matching models

Description: The authors introduce RealUID, a unified distillation framework that applies to multiple matching model families (diffusion, flow matching, bridge matching, stochastic interpolants). It unifies prior distillation methods (FGM, SiD, IBMD) under a single theoretical foundation using a linearization technique and connects them to inverse optimization.

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. Improved Distribution Matching Distillation for Fast Image Synthesis

URL: View paper

Brief Assessment

Improved Distribution Matching[32] focuses on distilling diffusion models using distribution matching without regression loss, but does not present a universal framework unifying multiple matching model families (diffusion, flow, bridge) under a single theoretical foundation as claimed in the original paper.

---

2. Multistep Distillation of Diffusion Models via Moment Matching

URL: View paper

Brief Assessment

Moment Matching Distillation[38] focuses on distilling diffusion models via moment matching for few-step sampling, not on unifying distillation methods across multiple matching model families (diffusion, flow, bridge) under a single framework with real data incorporation.

---

3. SparseFusion: Distilling View-Conditioned Diffusion for 3D Reconstruction

URL: View paper

Brief Assessment

SparseFusion[42] focuses on sparse-view 3D reconstruction by distilling view-conditioned diffusion models into 3D scene representations, not on universal distillation frameworks for matching models (diffusion, flow matching, bridge matching).

---

4. Adversarial Diffusion Distillation

URL: View paper

Brief Assessment

Adversarial Diffusion Distillation[40] focuses on distilling diffusion models using adversarial training combined with score distillation, not on unifying multiple matching model families (diffusion, flow matching, bridge matching, stochastic interpolants) under a single theoretical framework as RealUID does.

---

5. Simple and fast distillation of diffusion models

URL: View paper

Brief Assessment

Fast Diffusion Distillation[39] focuses on efficient distillation of diffusion models through trajectory matching with a teacher model, not on unifying multiple matching model families (diffusion, flow matching, bridge matching) under a single theoretical framework as RealUID does.

---

6. Score identity Distillation: Exponentially Fast Distillation of Pretrained Diffusion Models for One-Step Generation

URL: View paper

Brief Assessment

Score Identity Distillation[2] focuses specifically on distilling diffusion models using score-matching objectives and does not present a universal framework applicable to multiple matching model families (diffusion, flow matching, bridge matching, stochastic interpolants) as claimed in the original contribution.

7. One-Step Diffusion with Distribution Matching Distillation

URL: View paper

Brief Assessment

Distribution Matching Distillation[4] focuses on minimizing KL divergence for diffusion models specifically, not on unifying multiple matching model families (diffusion, flow matching, bridge matching, stochastic interpolants) under a single theoretical framework as RealUID does.

8. Diff-instruct: A universal approach for transferring knowledge from pre-trained diffusion models

URL: View paper

Brief Assessment

Diff-instruct[41] focuses on transferring knowledge from pre-trained diffusion models to other generative models (GANs, NERFs) using a novel IKL divergence, not on unifying distillation methods across diffusion, flow matching, and bridge matching frameworks as RealUID does.

9. Simple Distillation for One-Step Diffusion Models

URL: View paper

Brief Assessment

Simple Distillation[15] focuses on distilling multi-step diffusion models into one-step generators using contrastive energy distillation, not on unifying distillation methods across multiple matching model families (diffusion, flow matching, bridge matching, stochastic interpolants) as RealUID does.

10. On Distillation of Guided Diffusion Models

URL: View paper

Brief Assessment

Guided Diffusion Distillation[37] focuses specifically on distilling classifier-free guided diffusion models for faster sampling, not on creating a universal framework that unifies distillation across multiple matching model families (diffusion, flow matching, bridge matching, stochastic interpolants) as claimed in the original contribution.

Contribution 2: Natural incorporation of real data without GANs

Description: The authors propose a method to integrate real data into the distillation procedure by modifying the universal matching loss with weighting coefficients alpha and beta. This approach avoids the architectural modifications and optimization challenges associated with adversarial training and discriminator networks.

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. On the diversity and realism of distilled dataset: An efficient dataset distillation paradigm

URL: View paper

Brief Assessment

Diversity Realism Distillation[48] addresses dataset distillation for image synthesis, not diffusion model distillation. The candidate focuses on compressing training datasets by selecting and stitching image patches, which is fundamentally different from the original paper's approach of modifying matching losses in generative model distillation.

2. Datadam: Efficient dataset distillation with attention matching

URL: View paper

Brief Assessment

Datadam[44] focuses on dataset distillation through attention matching and feature distribution alignment, not on incorporating real data into distillation procedures. The paper does not address the specific challenge of integrating real data without adversarial training that the original contribution claims to solve.

3. Training data-efficient image transformers & distillation through attention

URL: View paper

Brief Assessment

DeiT[45] focuses on knowledge distillation for vision transformers using a distillation token, not on incorporating real data into generative model distillation procedures. The paper addresses image classification tasks with transformers, while the original contribution concerns generative matching models (diffusion/flow models) and their distillation framework.

4. Rethinking data distillation: Do not overlook calibration

URL: View paper

Brief Assessment

Calibration Data Distillation[51] addresses calibration of networks trained on distilled data through masking techniques, not the incorporation of real data into distillation procedures. The candidate focuses on calibration quality rather than real data integration methods.

5. Reliable data distillation on graph convolutional network

URL: View paper

Brief Assessment

Reliable Graph Distillation[52] focuses on graph convolutional networks and knowledge distillation for semi-supervised node classification, not on diffusion/flow model distillation or avoiding GANs in generative model training.

6. What is dataset distillation learning?

URL: View paper

Brief Assessment

What is Dataset Distillation[50] focuses on analyzing what information is captured in distilled datasets and how individual distilled data points contain semantic information. It does not address methods for incorporating real data into distillation procedures or propose alternatives to adversarial training approaches.

7. Dataset distillation by matching training trajectories

URL: View paper

Brief Assessment

Training Trajectory Matching[43] focuses on dataset distillation by matching training trajectories between synthetic and real data, but does not address the specific problem of incorporating real data into distillation procedures without adversarial losses or discriminators.

---

8. In-context data distillation with TabPFN

URL: View paper

Brief Assessment

In-context TabPFN[49] addresses tabular data classification through context optimization, not generative model distillation. It does not involve matching models, diffusion frameworks, or adversarial training architectures that the original contribution specifically targets.

---

9. Dataset distillation via the wasserstein metric

URL: View paper

Brief Assessment

Wasserstein Dataset Distillation[47] focuses on dataset distillation using optimal transport theory and Wasserstein metrics for distribution matching, not on incorporating real data into distillation procedures or avoiding adversarial training frameworks.

---

10. Unifying distillation and privileged information

URL: View paper

Brief Assessment

Unifying Distillation Privileged[46] focuses on distillation between different feature representations (privileged vs. regular features) for knowledge transfer, not on incorporating real data into distillation procedures to avoid adversarial training frameworks.

Contribution 3: Universal Matching loss with real data (RealUM)

Description: The authors define a new loss function that combines terms for both generated and real data distributions, parameterized by coefficients alpha and beta. This loss preserves the property that it yields the same teacher function when applied to real data, enabling real data incorporation while maintaining theoretical consistency.

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. Joint Velocity-Growth Flow Matching for Single-Cell Dynamics Modeling

URL: View paper

Brief Assessment

Velocity-Growth Flow Matching[57] addresses single-cell dynamics modeling with unpaired/unbalanced data via flow matching and optimal transport. It does not propose unified matching loss functions combining generated and real data distributions for distillation frameworks.

---

2. Improved Distribution Matching Distillation for Fast Image Synthesis

URL: View paper

Prior Art Analysis

Improved Distribution Matching[32] demonstrates prior work on incorporating real data into distillation losses without GANs. The candidate paper explicitly describes integrating a GAN objective where 'the discriminator is trained to distinguish between real images and images produced by our generator' and states 'trained using real data, the gan classifier does not suffer from the teacher's limitation.' This shows that real data incorporation in distillation was explored before the original paper's RealUM contribution, though through a different mechanism (GAN-based vs. the original's unified loss formulation).

Evidence

Evidence 1 - **Rationale**: This pair shows that Improved Distribution Matching[32] already incorporated real data into distillation training, though using a GAN-based approach rather than the original paper's GAN-free unified loss formulation. - **Original**: we present realuid, a universal distillation framework for all matching models that seamlessly incorporates real data into the distillation procedure without gans. - **Candidate**: we address this issue by incorporating an additional gan objective into our pipeline, where the discriminator is trained to distinguish between real images and images produced by our generator. trained using real data, the gan classifier does not suffer from the teacher's limitation, potentially all...

---

3. Real-fake: Effective training data synthesis through distribution matching

URL: View paper

Brief Assessment

Real-fake Synthesis[59] focuses on training data synthesis through distribution matching for supervised learning tasks, not on distillation of generative models. The paper addresses matching synthetic and real data distributions for classification, which is a fundamentally different problem from the original paper's unified matching loss for distilling diffusion/flow models.

---

4. A generalised novel loss function for computational fluid dynamics

URL: View paper

Brief Assessment

CFD Loss Function[55] focuses on computational fluid dynamics applications with gradient-weighted loss functions for spatial field generation, not on matching models or distillation frameworks for generative modeling.

---

5. Dataset Condensation with Distribution Matching

URL: View paper

Brief Assessment

Dataset Condensation Matching[53] focuses on dataset condensation through feature distribution matching in embedding spaces, not on unified matching loss functions for generative model distillation combining real and generated data distributions as in the original paper.

---

6. Discriminative fault diagnosis transfer learning network under joint mechanism

URL: View paper

Brief Assessment

Discriminative Fault Diagnosis[58] focuses on fault diagnosis transfer learning using domain adaptation techniques (CORAL, JMMD) for mechanical systems, not on unified matching loss functions for generative models combining generated and real data distributions.

7. One-Step Diffusion with Distribution Matching Distillation

URL: View paper

Brief Assessment

Distribution Matching Distillation[4] uses KL divergence minimization with score function differences, not a unified matching loss combining generated and real data distributions with alpha/beta coefficients as described in RealUM.

8. Few-shot LLM Synthetic Data with Distribution Matching

URL: View paper

Brief Assessment

Few-shot LLM Synthesis[54] addresses synthetic data generation for LLMs using distribution matching via maximum mean discrepancy, not unified loss functions for distillation of generative models combining generated and real data distributions as in the original paper.

9. Imagebinddc: Compressing multi-modal data with imagebind-based condensation

URL: View paper

Brief Assessment

ImageBindDC[60] focuses on multi-modal data condensation using characteristic function discrepancy in a unified embedding space, not on unified matching loss functions combining generated and real data distributions for distillation of matching models.

10. A Unifying Generator Loss Function for Generative Adversarial Networks

URL: View paper

Brief Assessment

Unifying Generator Loss[56] focuses on GAN generator loss functions and Jensen-fα-divergence minimization, not on matching model distillation or real data incorporation in diffusion/flow frameworks.

Appendix: Text Similarity Detection

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

References

- [0] Universal Inverse Distillation for Matching Models with Real-Data Supervision (No GANs) View paper
- [1] Di o: Distilling masked diffusion models into one-step generator View paper
- [2] Score identity Distillation: Exponentially Fast Distillation of Pretrained Diffusion Models for One-Step Generation View paper
- [3] Adversarial distribution matching for diffusion distillation towards efficient image and video synthesis View paper
- [4] One-Step Diffusion with Distribution Matching Distillation View paper
- [5] Multi-student Diffusion Distillation for Better One-step Generators View paper
- [6] Self-Corrected Flow Distillation for Consistent One-Step and Few-Step Image Generation View paper
- [7] VideoScene: Distilling Video Diffusion Model to Generate 3D Scenes in One Step View paper
- [8] Sana-sprint: One-step diffusion with continuous-time consistency distillation View paper
- [9] Distilled Decoding 1: One-step Sampling of Image Auto-regressive Models with Flow Matching View paper
- [10] Learning Few-Step Diffusion Models by Trajectory Distribution Matching View paper
- [11] SwiftBrush: One-Step Text-to-Image Diffusion Model with Variational Score Distillation View paper
- [12] Inverse Bridge Matching Distillation View paper
- [13] One diffusion step to real-world super-resolution via flow trajectory distillation View paper
- [14] Few-step diffusion via score identity distillation View paper
- [15] Simple Distillation for One-Step Diffusion Models View paper
- [16] Towards One-step Causal Video Generation via Adversarial Self-Distillation View paper
- [17] Distilling Diffusion Models into Conditional GANs View paper
- [18] Diffusion Models Are Innate One-Step Generators View paper
- [19] One-Step Diffusion Distillation through Score Implicit Matching View paper
- [20] Adversarial score identity distillation: Rapidly surpassing the teacher in one step View paper
- [21] Flow Generator Matching View paper
- [22] Revisiting Diffusion Models: From Generative Pre-training to One-Step Generation View paper
- [23] Flow2GAN: Hybrid Flow Matching and GAN with Multi-Resolution Network for Few-step High-Fidelity Audio Generation View paper
- [24] Phased DMD: Few-step Distribution Matching Distillation via Score Matching within Subintervals View paper
- [25] Regularized Distribution Matching Distillation for One-step Unpaired Image-to-Image Translation View paper
- [26] One-step Diffusion Models with F-Divergence Distribution Matching View paper
- [27] VarDiU: A Variational Diffusive Upper Bound for One-Step Diffusion Distillation View paper
- [28] MeanFlow Transformers with Representation Autoencoders View paper
- [29] Guided Score identity Distillation for Data-Free One-Step Text-to-Image Generation View paper
- [30] ProtFlow: Fast Protein Sequence Design via Flow Matching on Compressed Protein Language Model Embeddings View paper
- [31] One-step Diffusion Models with Bregman Density Ratio Matching View paper
- [32] Improved Distribution Matching Distillation for Fast Image Synthesis View paper
- [33] EM Distillation for One-step Diffusion Models View paper
- [34] Robust One-step Speech Enhancement via Consistency Distillation View paper
- [35] One-step Flow Matching Generators View paper
- [36] On Distilling Generator Matching Models View paper
- [37] On Distillation of Guided Diffusion Models View paper
- [38] Multistep Distillation of Diffusion Models via Moment Matching View paper
- [39] Simple and fast distillation of diffusion models View paper
- [40] Adversarial Diffusion Distillation View paper

- [41] Diff-instruct: A universal approach for transferring knowledge from pre-trained diffusion models View paper
- [42] SparseFusion: Distilling View-Conditioned Diffusion for 3D Reconstruction View paper
- [43] Dataset distillation by matching training trajectories View paper
- [44] Datadam: Efficient dataset distillation with attention matching View paper
- [45] Training data-efficient image transformers & distillation through attention View paper
- [46] Unifying distillation and privileged information View paper
- [47] Dataset distillation via the wasserstein metric View paper
- [48] On the diversity and realism of distilled dataset: An efficient dataset distillation paradigm View paper
- [49] In-context data distillation with TabPFN View paper
- [50] What is dataset distillation learning? View paper
- [51] Rethinking data distillation: Do not overlook calibration View paper
- [52] Reliable data distillation on graph convolutional network View paper
- [53] Dataset Condensation with Distribution Matching View paper
- [54] Few-shot LLM Synthetic Data with Distribution Matching View paper
- [55] A generalised novel loss function for computational fluid dynamics View paper
- [56] A Unifying Generator Loss Function for Generative Adversarial Networks View paper
- [57] Joint Velocity-Growth Flow Matching for Single-Cell Dynamics Modeling View paper
- [58] Discriminative fault diagnosis transfer learning network under joint mechanism View paper
- [59] Real-fake: Effective training data synthesis through distribution matching View paper
- [60] Imagebinddc: Compressing multi-modal data with imagebind-based condensation View paper