Novelty Assessment Report

Paper: Efficient Zero-shot Inpainting with Decoupled Diffusion Guidance
PDF URL: https://openreview.net/pdf?id=5F93RfQ12T
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2026-01-07

Abstract

Diffusion models have emerged as powerful priors for image editing tasks such as inpainting and local modification, where the objective is to generate realistic content that remains consistent with observed regions. In particular, zero-shot approaches that leverage a pretrained diffusion model, without any retraining, have been shown to achieve highly effective reconstructions. However, state-of-the-art zero-shot methods typically rely on a sequence of surrogate likelihood functions, whose scores are used as proxies for the ideal score. This procedure however requires vector-Jacobian products through the denoiser at every reverse step, introducing significant memory and runtime overhead. To address this issue, we propose a new likelihood surrogate that yields simple and efficient to sample Gaussian posterior transitions, sidestepping the backpropagation through the denoiser network. Our extensive experiments show that our method achieves strong observation consistency compared with fine-tuned baselines and produces coherent, high-quality reconstructions, all while significantly reducing inference cost.

Core Task Landscape

This paper addresses: **Zero-Shot Image Inpainting Using Pretrained Diffusion Models**

A total of **36 papers** were analyzed and organized into a taxonomy with **17 categories**.

Taxonomy Overview

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

- **Diffusion Model Adaptation Mechanisms**
- **Multimodal and Conditional Inpainting**
- **Training-Based and Hybrid Approaches**
- **Domain-Specific Inpainting Applications**
- **Diffusion Model Enhancements and Efficiency**

Complete Taxonomy Tree

- Zero-Shot Image Inpainting Using Pretrained Diffusion Models Survey Taxonomy
- Diffusion Model Adaptation Mechanisms
  - Null-Space and Range-Space Guidance ★ (3 papers)
  - [0] Efficient Zero-shot Inpainting with Decoupled Diffusion Guidance (Anon et al., 2026) View paper
  - [1] Zero-Shot Image Inpainting using Pretrained Latent Diffusion Models (Yusuke Kakinuma, 2025) View paper
  - [2] Zero-Shot Image Restoration Using Denoising Diffusion Null-Space Model (Wang Yinhuai, 2022) View paper
  - Gradient and Attention Guidance (3 papers)
  - [11] GradPaint: Gradient-guided inpainting with diffusion models (Asya Grechka, 2024) View paper
  - [12] Magicremover: Tuning-free text-guided image inpainting with diffusion models (Yang Si-yuan, 2023) View paper
  - [23] GuidPaint: Class-Guided Image Inpainting with Diffusion Models (Wang Qi-min, 2025) View paper
  - Latent Space Optimization and Regularization (3 papers)
  - [7] Latentpaint: Image inpainting in latent space with diffusion models (Ciprian Corneanu, 2024) View paper
  - [16] Coherent and Multi-modality Image Inpainting via Latent Space Optimization (Pan Ling-zhi, 2024) View paper
  - [26] Exploiting Generative Diffusion Prior With Latent Low-Rank Regularization for Image Inpainting (Zhentao Zou, 2024) View paper
  - Stochastic Sampling and Resampling Strategies (3 papers)
  - [8] Zero-shot adaptation for approximate posterior sampling of diffusion models in inverse problems (AkÃ§akaya, 2024) View paper
  - [19] A Latent Space of Stochastic Diffusion Models for Zero-Shot Image Editing and Guidance (Chen Henry Wu, 2023) View paper
  - [33] RePaint: Inpainting using Denoising Diffusion Probabilistic Models (Andreas Lugmayr, 2022) View paper
- Multimodal and Conditional Inpainting
  - Text and Prompt-Guided Inpainting (2 papers)
  - [3] Uni-paint: A unified framework for multimodal image inpainting with pretrained diffusion model (Yang Shi-yuan, 2023) View paper
  - [10] Hd-painter: high-resolution and prompt-faithful text-guided image inpainting with diffusion models (Manukyan, 2023) View paper
  - Exemplar and Subject-Driven Inpainting (2 papers)
  - [14] Freecompose: Generic zero-shot image composition with diffusion prior (Zhekai Chen, 2024) View paper
  - [17] Paste, inpaint and harmonize via denoising: Subject-driven image editing with pre-trained diffusion model (Zhang Xin, 2023) View paper
  - Sketch and Spatial Guidance (2 papers)

    - [13] DesignEdit: Unify Spatial-Aware Image Editing via Training-free Inpainting with a Multi-Layered Latent Diffusion Framework (Jia, 2025) View paper
    - [18] Zero-Shot Depth Aware Image Editing with Diffusion Models (R Parihar, 2025) View paper
    - Multi-Modality Fusion Frameworks (1 papers)
    - [32] MaGIC: Multi-modality Guided Image Completion (Yu, 2023) View paper
- Training-Based and Hybrid Approaches
    - Supervised and Fine-Tuned Inpainting Models (1 papers)
    - [36] Image Inpainting via Jointing Structure Restoration and End-to-end Reversible Diffusion (S Li, n.d.) View paper
    - Hybrid Supervised and Zero-Shot Frameworks (1 papers)
    - [15] InvFussion: Bridging Supervised and Zero-shot Diffusion for Inverse Problems (Elata, 2025) View paper
- Domain-Specific Inpainting Applications
    - Medical Image Inpainting and Restoration (3 papers)
    - [5] Zero-Shot Medical Image Translation via Frequency-Guided Diffusion Models (Yunxiang Li, 2023) View paper
    - [28] Inpainting is All You Need: A Diffusion-based Augmentation Method for Semi-supervised Medical Image Segmentation (Hu Xinrong, 2025) View paper
    - [31] Unsupervised anomaly localization in high-resolution breast scans using deep pluralistic image completion (Nicholas Konz, 2023) View paper
    - Document and Historical Artifact Restoration (1 papers)
    - [9] An attempt at zero-shot ancient documents restoration based on diffusion models (Hayata Kaneko, 2023) View paper
    - Video and Temporal Inpainting (3 papers)
    - [6] OmnimatteZero: Fast Training-free Omnimatte with Pre-trained Video Diffusion Models (Samuel, 2025) View paper
    - [35] ZeroPatcher: Training-free Sampler for Video Inpainting and Editing (S Yang, n.d.) View paper
    - Specialized Imaging Modalities (4 papers)
    - [20] Omnissr: Zero-shot omnidirectional image super-resolution using stable diffusion model (Runyi Li, 2024) View paper
    - [24] Language-guided manipulation with diffusion policies and constrained inpainting (Hao Ce, 2024) View paper
    - [29] Facial Landmark Detection and Head Pose Estimation of Occluded Faces Based on Stable Diffusion Image Completion (Wenzhangzhi Guo, 2025) View paper
    - [30] Generate Realistic Image and Segmentation Pair using Diffusion Probabilistic Model (Zhou, 2023) View paper
- Diffusion Model Enhancements and Efficiency
    - High-Resolution and Upsampling Techniques (1 papers)
    - [4] Upsample Guidance: Scale Up Diffusion Models without Training (Hwang, 2024) View paper
    - Fast Sampling and Consistency Models (2 papers)
    - [25] Consistency Models (Song Yang, 2023) View paper
    - [34] LanPaint: Training-Free Diffusion Inpainting with Asymptotically Exact and Fast Conditional Sampling (Zheng, n.d.) View paper
    - Generalization and Robustness (2 papers)
    - [22] Moderating the Generalization of Score-based Generative Model (Jiang Wan, 2024) View paper
    - [27] Zero-shot Depth Completion via Test-time Alignment with Affine-invariant Depth Prior (Byung Ki Kwon, 2025) View paper

Narrative

Core task: Zero-shot image inpainting using pretrained diffusion models. The field has organized itself around several complementary directions. Diffusion Model Adaptation Mechanisms explores how to steer pretrained models without retraining, often through guidance strategies that manipulate the denoising process to respect masked regions while generating coherent content. Multimodal and Conditional Inpainting extends these ideas by incorporating text prompts, depth maps, or other modalities to control what gets synthesized. Training-Based and Hybrid Approaches blend zero-shot flexibility with lightweight fine-tuning or test-time optimization to improve quality or domain fit. Domain-Specific Inpainting Applications targets specialized use cases such as face completion, document restoration, or video layer matting, while Diffusion Model Enhancements and Efficiency focuses on accelerating sampling or reducing computational overhead. Together, these branches reflect a tension between leveraging off-the-shelf generative priors and adapting them to diverse constraints and modalities.

Within Diffusion Model Adaptation Mechanisms, a particularly active line of work centers on null-space and range-space guidance, where methods like Null-Space Model[2] and Pretrained Latent Inpainting[1] decompose the generation process to preserve known pixels while freely hallucinating missing content. Decoupled Diffusion Guidance[0] sits squarely in this cluster, proposing a refined decomposition that separates constraints from creative synthesis more cleanly than earlier approaches. Compared to Null-Space Model[2], which introduced the foundational projection idea, Decoupled Diffusion Guidance[0] emphasizes decoupling guidance signals to reduce artifacts at mask boundaries. Meanwhile, works like Pretrained Latent Inpainting[1] operate in latent space for efficiency, raising questions about how best to balance pixel-level fidelity with computational cost. Across these studies, the central challenge remains achieving seamless blending and semantic coherence without task-specific training, a goal that continues to drive innovation in guidance design and sampling strategies.

Related Works in Same Category

The following **2 sibling papers** share the same taxonomy leaf node with the original paper:

1. Zero-Shot Image Inpainting using Pretrained Latent Diffusion Models

Authors: Yusuke Kakinuma, Takamichi Miyata, Kaito Hosono, Hirotsugu Kinoshita | **Year/Venue**: 2025 | **URL**: View paper

Abstract

The DDNM (Denoising Diffusion Null-space Model), a zero-shot image restoration method, uses a pre-trained diffusion-based image generation model and can be applied to a variety of image restoration tasks without task-specific training. However, since DDNM uses a diffusion model trained by the ImageNet as its backbone, its restoration capability is strongly constrained by the limited number of classes of training images. On the other hand, latent diffusion models (LDMs), which perform diffusion-b...

Relationship Analysis

Both papers belong to the Null-Space and Range-Space Guidance category, leveraging pretrained diffusion models for zero-shot inpainting by constraining the reverse diffusion process with degradation operators. They overlap in their zero-shot approach and use of range-space projections to maintain consistency with observed regions. However, the original paper focuses on efficient Gaussian posterior transitions that avoid backpropagation through the denoiser, while the candidate paper specifically adapts the DDNM framework to latent diffusion models (LDMs) by addressing the challenge of nonlinear encoders in latent space while preserving spatial features for inpainting.

2. Zero-Shot Image Restoration Using Denoising Diffusion Null-Space Model

Authors: Wang Yinhuai, Yu Jiwen, Yinhuai Wang, Zhang Jian, Jiwen Yu, et al. (6 authors total) | **Year/Venue**: 2022 | **URL**: View paper

Abstract

Most existing Image Restoration (IR) models are task-specific, which can not be generalized to different degradation operators. In this work, we propose the Denoising Diffusion Null-Space Model (DDNM), a novel zero-shot framework for arbitrary linear IR problems, including but not limited to image super-resolution, colorization, inpainting, compressed sensing, and deblurring. DDNM only needs a pre-trained off-the-shelf diffusion model as the generative prior, without any extra training or networ...

Relationship Analysis

Both papers belong to the Null-Space and Range-Space Guidance category, using range-null space decomposition to constrain diffusion reverse processes for zero-shot image restoration. They overlap in applying null-space projections to ensure data consistency while leveraging pretrained diffusion models without retraining. The key difference is that DDNM iteratively refines only null-space contents during reverse diffusion with explicit range-null decomposition ($A^\dagger y + (I - A^\dagger A)x_0$), while the original paper proposes a new likelihood surrogate yielding Gaussian posterior transitions that sidesteps backpropagation through the denoiser, focusing on computational efficiency rather than explicit null-space manipulation.

Contributions Analysis

Overall novelty summary. The paper proposes a vector-Jacobian-product-free framework for zero-shot diffusion-based inpainting, introducing a new likelihood surrogate that yields Gaussian posterior transitions without backpropagation through the denoiser. It resides in the Null-Space and Range-Space Guidance leaf, which contains only three papers total. This leaf sits within the broader Diffusion Model Adaptation Mechanisms branch, indicating a moderately crowded research direction focused on steering pretrained models without retraining. The small leaf size suggests this specific projection-based guidance approach represents a focused subfield rather than a saturated research area.

The taxonomy reveals that neighboring leaves explore alternative guidance mechanisms: Gradient and Attention Guidance manipulates sampling through optimization or attention, Latent Space Optimization regularizes representations during diffusion, and Stochastic Sampling modifies noise schedules. The paper's null-space approach differs fundamentally by decomposing the generation process to preserve observed pixels while hallucinating missing content, contrasting with gradient-based methods that iteratively refine outputs. This positioning suggests the work builds on a distinct lineage of projection-based techniques rather than gradient or attention manipulation, though all share the zero-shot adaptation goal.

Among nineteen candidates examined, the VJP-free framework contribution shows one refutable candidate from five examined, indicating some prior work addresses computational efficiency in zero-shot inpainting. The decoupled twisting function examined four candidates with none refutable, suggesting this theoretical formulation may be more novel. The DING method examined ten candidates without refutation, though this larger search scope does not guarantee exhaustive coverage. The limited search scale means these statistics reflect top-semantic-match overlap rather than comprehensive field assessment, leaving open whether deeper literature contains additional relevant work.

Based on the constrained search of nineteen papers, the work appears to occupy a moderately explored niche within zero-shot diffusion adaptation. The efficiency focus and theoretical decomposition show partial novelty, though the single refutable pair for the core framework suggests some computational concerns have been addressed previously. The analysis covers top-semantic matches and citation expansion but does not claim exhaustive field coverage, particularly for recent preprints or domain-specific efficiency techniques outside the main inpainting literature.

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

Contribution 1: VJP-free framework for zero-shot inpainting with diffusion priors

Description: The authors introduce a framework that eliminates the need for vector-Jacobian product evaluations and backpropagation through the denoiser network, addressing the computational and memory overhead of existing zero-shot methods.

This contribution was assessed against **5 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Fast constrained sampling in pre-trained diffusion models

URL: View paper

Prior Art Analysis

Fast Constrained Sampling[41] demonstrates that prior work exists which eliminates the need for vector-Jacobian product evaluations and backpropagation through the denoiser network for zero-shot image generation tasks. The candidate paper explicitly addresses the same computational overhead problem and proposes an alternative method that avoids backpropagation through the denoiser weights, requiring only forward passes. Both papers identify the same limitation in existing methods (computational cost of backpropagation through the denoiser) and propose solutions that eliminate this requirement, though using different technical approaches.

Evidence

Evidence 1 - **Rationale**: Both papers identify backpropagation through the denoiser as a key limitation and propose methods to eliminate it for zero-shot generation tasks. - **Original**: we propose a new vjp-free framework for zero-shot inpainting with a pre-trained diffusion prior. our key idea is to approximate the intractable twisted posterior-sampling transitions by a closed-form mixture distribution that can be sampled exactly, thereby eliminating the need for vjp evaluations a... - **Candidate**: previous approaches either utilized backpropagation through the denoiser network, making them significantly slower and more memory-demanding than simple text-to-image generation, or only enforced the constraint locally, failing to capture critical long-range correlations in the sampled image. in thi...

Evidence 2 - **Rationale**: Both papers explicitly address the computational cost of backpropagation and VJP evaluations through the denoiser, and propose methods requiring only forward passes instead. - **Original**: while current zero-shot methods are appealing, they face a critical practical limitation. implementations of strong zero-shot posterior sampling with diffusion priors typically rely on the twisting function proposed by ho et al. (2022); chung et al. (2023); song et al. (2023a), which corresponds to ... - **Candidate**: the direction of optimization we propose in eq. (10) has another advantage over the gradient descent update. the direction et can be computed numerically to save both on computation and memory compared to using backpropagation on the cost in eq. (1). to derive the update, consider the function h(s) ...

Evidence 3 - **Rationale**: Both papers demonstrate that their VJP-free approaches significantly reduce memory consumption compared to methods requiring backpropagation through the denoiser. - **Original**: this makes such methods computationally demanding, memory intensive, and often slower than training a dedicated conditional model. contributions. we propose a new vjp-free framework for zero-shot inpainting with a pre-trained diffusion prior. - **Candidate**: regarding memory, for a single image, our forward passes only require~9gb of memory, compared to backpropagation, which consumes~17gb.

---

2. Training-Free Safe Denoisers for Safe Use of Diffusion Models

URL: View paper

Brief Assessment

Safe Denoisers[43] addresses a different problem domain (safety in diffusion models) rather than zero-shot inpainting efficiency. The candidate focuses on modifying sampling trajectories to avoid unsafe content, not on eliminating VJP computations for inpainting tasks.

---

3. DreamShot: Teaching Cinema Shots to Latent Diffusion Models

URL: View paper

Brief Assessment

DreamShot[42] focuses on teaching cinema shots to latent diffusion models for cinematic content generation, not on zero-shot inpainting or eliminating vector-Jacobian product evaluations for inverse problems.

---

4. LanPaint: Training-Free Diffusion Inpainting with Asymptotically Exact and Fast Conditional Sampling

URL: View paper

Brief Assessment

LanPaint[34] addresses zero-shot inpainting but uses a fundamentally different approach based on Langevin dynamics Monte Carlo rather than VJP-free likelihood surrogates. The candidate focuses on bidirectional guided scores and fast Langevin dynamics, not on eliminating vector-Jacobian products through decoupled likelihood approximations.

---

5. Zero-Shot Solving of Imaging Inverse Problems via Noise-Refined Likelihood Guided Diffusion Models

URL: View paper

Brief Assessment

Noise-Refined Likelihood[44] focuses on general linear inverse problems (super-resolution, deblurring, compressive sensing) using noise refinement in the prediction space, not specifically on inpainting tasks. The technical approach differs fundamentally from the original paper's decoupled guidance mechanism for inpainting.

Contribution 2: Decoupled twisting function with closed-form mixture distribution

Description: The method modifies the twisting function by evaluating the denoiser at an independent draw from the pretrained transition, breaking the dependency and enabling exact sampling from posterior transitions without VJP computations.

This contribution was assessed against **4 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Particle denoising diffusion sampler

URL: View paper

Brief Assessment

Particle Denoising[39] focuses on Monte Carlo sampling from unnormalized densities using particle methods, not on zero-shot inpainting with diffusion models. The technical approaches differ fundamentally in application domain and methodology.

---

2. A Mixture-Based Framework for Guiding Diffusion Models

URL: View paper

Brief Assessment

Mixture-Based Framework[38] uses mixture approximations of intermediate posterior distributions with Gibbs sampling, while the original paper proposes decoupling the denoiser evaluation from transition density arguments to enable exact Gaussian posterior sampling without VJP computations. These are distinct technical approaches to posterior sampling in diffusion models.

---

3. Divide-and-conquer posterior sampling for denoising diffusion priors

URL: View paper

Brief Assessment

Divide-and-conquer Sampling[37] focuses on approximating posterior transitions through variational inference and bridge-kernel smoothing in a divide-and-conquer framework, not on decoupling the twisting function to enable closed-form mixture distributions without VJP computations as in the original paper.

---

4. Diffusion Bridge Mixture Transports, SchrÃ¶dinger Bridge Problems and Generative Modeling

URL: View paper

Brief Assessment

Diffusion Bridge Mixture[40] addresses Schrödinger bridge problems and optimal transport between probability measures using diffusion processes, not posterior sampling transitions in diffusion-based image editing. The technical focus and application domain differ fundamentally from the original paper's zero-shot inpainting framework.

Contribution 3: DING method for efficient zero-shot inpainting

Description: The authors develop DING, which achieves superior trade-offs between fidelity and realism while being faster and more memory-efficient than competing approaches, even outperforming fine-tuned models without task-specific training.

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. Latino-pro: Latent consistency inverse solver with prompt optimization

URL: View paper

Brief Assessment

Latino-pro[47] focuses on inverse problems using latent consistency models with prompt optimization, not specifically on zero-shot inpainting with decoupled diffusion guidance. The candidate addresses different technical challenges (prompt calibration, consistency models) compared to DING's decoupled guidance approach.

---

2. Torch-advent-civilization-evolution: accelerating diffusion model for image restoration

URL: View paper

Brief Assessment

Torch-advent[51] focuses on image restoration tasks (super-resolution, deblurring, colorization) using a 'torch' initialization strategy, not on inpainting with decoupled guidance. The technical approaches differ fundamentally in their methodology and target applications.

---

3. Zero-Shot Depth Aware Image Editing with Diffusion Models

URL: View paper

Brief Assessment

Depth Aware Editing[18] focuses on depth-aware image editing tasks (object insertion and scene compositing at specified depths) using layered representations, not on efficient zero-shot inpainting methods optimizing fidelity-realism tradeoffs under low NFE budgets.

---

4. Conditional neural field latent diffusion model for generating spatiotemporal turbulence

URL: View paper

Brief Assessment

Spatiotemporal Turbulence[45] focuses on generating spatiotemporal turbulent flows in fluid dynamics using conditional neural field latent diffusion models, not on zero-shot image inpainting methods. The domains and technical approaches are fundamentally different.

---

5. D-flow: Differentiating through flows for controlled generation

URL: View paper

Brief Assessment

D-flow[46] focuses on controlled generation via differentiating through flow models for inverse problems, while DING addresses efficient zero-shot inpainting by decoupling diffusion guidance to avoid VJP computations. The technical approaches and optimization strategies differ fundamentally.

---

6. Act-diffusion: Efficient adversarial consistency training for one-step diffusion models

URL: View paper

Brief Assessment

Act-diffusion[48] focuses on accelerating diffusion models through adversarial consistency training for one-step generation, not on zero-shot inpainting methods or achieving fidelity-realism trade-offs under low NFE budgets.

---

7. Zero-Shot Image Restoration via Few-Step Guidance of Consistency Models

URL: View paper

Brief Assessment

Few-Step Guidance[52] focuses on consistency models (CMs) for image restoration with 4 NFEs, while DING uses diffusion models with decoupled guidance achieving superior performance at 50 NFEs. The technical approaches and model architectures differ fundamentally.

---

8. Masked pre-training enables universal zero-shot denoiser

URL: View paper

Brief Assessment

Universal Denoiser[49] focuses on zero-shot image denoising using masked pre-training on natural images, not zero-shot inpainting. The candidate addresses noise removal across various noise types (Gaussian, S&P, speckle, etc.), while DING targets inpainting tasks with missing pixel regions. These are fundamentally different inverse problems with distinct technical approaches.

---

9. Dreamsampler: Unifying diffusion sampling and score distillation for image manipulation

URL: View paper

Brief Assessment

DreamSampler[50] focuses on unifying diffusion sampling and score distillation for image manipulation tasks including editing and vectorization, not specifically on efficient zero-shot inpainting with decoupled guidance achieving superior fidelity-realism trade-offs under low NFE budgets.

---

10. Automated mural restoration via semi supervised segmentation and prompt guided diffusion inpainting

URL: View paper

Brief Assessment

Mural Restoration[53] focuses on cultural heritage preservation through semi-supervised segmentation and prompt-guided diffusion inpainting for murals, not general-purpose zero-shot inpainting methods achieving high fidelity under low NFE budgets.

Appendix: Text Similarity Detection

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

References

- [0] Efficient Zero-shot Inpainting with Decoupled Diffusion Guidance View paper
- [1] Zero-Shot Image Inpainting using Pretrained Latent Diffusion Models View paper
- [2] Zero-Shot Image Restoration Using Denoising Diffusion Null-Space Model View paper
- [3] Uni-paint: A unified framework for multimodal image inpainting with pretrained diffusion model View paper
- [4] Upsample Guidance: Scale Up Diffusion Models without Training View paper
- [5] Zero-Shot Medical Image Translation via Frequency-Guided Diffusion Models View paper
- [6] OmnimatteZero: Fast Training-free Omnimatte with Pre-trained Video Diffusion Models View paper
- [7] Latentpaint: Image inpainting in latent space with diffusion models View paper
- [8] Zero-shot adaptation for approximate posterior sampling of diffusion models in inverse problems View paper
- [9] An attempt at zero-shot ancient documents restoration based on diffusion models View paper
- [10] Hd-painter: high-resolution and prompt-faithful text-guided image inpainting with diffusion models View paper
- [11] GradPaint: Gradient-guided inpainting with diffusion models View paper
- [12] Magicremover: Tuning-free text-guided image inpainting with diffusion models View paper
- [13] DesignEdit: Unify Spatial-Aware Image Editing via Training-free Inpainting with a Multi-Layered Latent Diffusion Framework View paper
- [14] Freecompose: Generic zero-shot image composition with diffusion prior View paper
- [15] InvFussion: Bridging Supervised and Zero-shot Diffusion for Inverse Problems View paper
- [16] Coherent and Multi-modality Image Inpainting via Latent Space Optimization View paper
- [17] Paste, inpaint and harmonize via denoising: Subject-driven image editing with pre-trained diffusion model View paper
- [18] Zero-Shot Depth Aware Image Editing with Diffusion Models View paper

- [19] A Latent Space of Stochastic Diffusion Models for Zero-Shot Image Editing and Guidance View paper
- [20] Omnissr: Zero-shot omnidirectional image super-resolution using stable diffusion model View paper
- [21] OmnimatteZero: Training-free Real-time Omnimatte with Pre-trained Video Diffusion Models View paper
- [22] Moderating the Generalization of Score-based Generative Model View paper
- [23] GuidPaint: Class-Guided Image Inpainting with Diffusion Models View paper
- [24] Language-guided manipulation with diffusion policies and constrained inpainting View paper
- [25] Consistency Models View paper
- [26] Exploiting Generative Diffusion Prior With Latent Low-Rank Regularization for Image Inpainting View paper
- [27] Zero-shot Depth Completion via Test-time Alignment with Affine-invariant Depth Prior View paper
- [28] Inpainting is All You Need: A Diffusion-based Augmentation Method for Semi-supervised Medical Image Segmentation View paper
- [29] Facial Landmark Detection and Head Pose Estimation of Occluded Faces Based on Stable Diffusion Image Completion View paper
- [30] Generate Realistic Image and Segmentation Pair using Diffusion Probabilistic Model View paper
- [31] Unsupervised anomaly localization in high-resolution breast scans using deep pluralistic image completion View paper
- [32] MaGIC: Multi-modality Guided Image Completion View paper
- [33] RePaint: Inpainting using Denoising Diffusion Probabilistic Models View paper
- [34] LanPaint: Training-Free Diffusion Inpainting with Asymptotically Exact and Fast Conditional Sampling View paper
- [35] ZeroPatcher: Training-free Sampler for Video Inpainting and Editing View paper
- [36] Image Inpainting via Jointing Structure Restoration and End-to-end Reversible Diffusion View paper
- [37] Divide-and-conquer posterior sampling for denoising diffusion priors View paper
- [38] A Mixture-Based Framework for Guiding Diffusion Models View paper
- [39] Particle denoising diffusion sampler View paper
- [40] Diffusion Bridge Mixture Transports, SchrÃ¶dinger Bridge Problems and Generative Modeling View paper
- [41] Fast constrained sampling in pre-trained diffusion models View paper
- [42] DreamShot: Teaching Cinema Shots to Latent Diffusion Models View paper
- [43] Training-Free Safe Denoisers for Safe Use of Diffusion Models View paper
- [44] Zero-Shot Solving of Imaging Inverse Problems via Noise-Refined Likelihood Guided Diffusion Models View paper
- [45] Conditional neural field latent diffusion model for generating spatiotemporal turbulence View paper
- [46] D-flow: Differentiating through flows for controlled generation View paper
- [47] Latino-pro: Latent consistency inverse solver with prompt optimization View paper
- [48] Act-diffusion: Efficient adversarial consistency training for one-step diffusion models View paper
- [49] Masked pre-training enables universal zero-shot denoiser View paper
- [50] Dreamsampler: Unifying diffusion sampling and score distillation for image manipulation View paper
- [51] Torch-advent-civilization-evolution: accelerating diffusion model for image restoration View paper
- [52] Zero-Shot Image Restoration via Few-Step Guidance of Consistency Models View paper
- [53] Automated mural restoration via semi supervised segmentation and prompt guided diffusion inpainting View paper