Project Proposal: Inverse Imaging via Drifting Generative Models

Zixuan Zhao

Motivation

Inverse imaging problems such as **noisy deconvolution** and **super-resolution** are classic ill-posed problems, where multiple clean images can explain the same degraded observation. Recent progress shows that modern generative models can act as strong image priors for solving such tasks, especially diffusion-based methods and related score-based approaches. However, many of these methods require iterative inference, which can be computationally expensive.

This project proposes to investigate whether **Drifting Generative Models** (DGM) [1] can be used as efficient priors for inverse imaging. Drifting models are especially attractive because they naturally support **one-step (1-NFE) inference**, potentially enabling much faster reconstruction than iterative samplers. We will study whether this one-step generative paradigm can still produce useful priors for restoration tasks, and we will compare **pixel-space** and **latent-space** variants (motivated by recent representation autoencoder work [2]).

Related Work

Our project sits at the intersection of generative modeling and inverse problems.

- **Diffusion models for inverse problems.** Diffusion and score-based models have become a major framework for image restoration and inverse problem solving, due to their ability to model rich natural image priors. A recent survey provides a broad taxonomy and evaluation of diffusion-based inverse problem methods [3].
- **Drifting Generative Models.** Deng et al. introduce a new generative paradigm in which the pushforward distribution evolves during training and supports one-step generation at inference time [1]. This suggests a promising direction for computationally efficient inverse reconstruction.
- **Latent representations / representation autoencoders.** Recent work on representation autoencoders (RAEs) for diffusion transformers argues that stronger latent spaces can improve generation quality and scalability [2]. This motivates our comparison of pixel-space versus latent-space drifting models for inverse tasks.

Project Overview and Final Goals

We will evaluate drifting-based generative priors on two inverse imaging tasks: Noisy deconvolution and super-resolution. To keep iteration fast and experiments manageable, we will use **MNIST** and **CIFAR-10** as primary evaluation datasets.

Core questions.

- Can a drifting generative model be adapted to produce good reconstructions for noisy deconvolution and super-resolution?

- How does reconstruction quality compare between **pixel-space** and **latent-space** drifting formulations?
- What are the speed/quality tradeoffs relative to simple baselines (e.g., bicubic upsampling / Wiener-style deconvolution / direct regression baseline)?

Expected final deliverables.

- A working drifting-based inverse reconstruction pipeline for at least one task (target: both tasks).
- Quantitative evaluation on MNIST and CIFAR-10 (e.g., PSNR, SSIM, MSE).
- Qualitative examples showing success/failure cases.
- A comparison of pixel-space vs. latent-space drifting (full or partial, depending on implementation complexity).

Milestones and Timeline

- **Week 1:** Read core papers; finalize task definitions, degradation models, and metrics; build data preprocessing + evaluation pipeline for MNIST/CIFAR-10.
- **Week 2:** Implement baseline inverse solvers and drifting-based reconstruction pipeline for one task (likely super-resolution first). Extend to noisy deconvolution; run initial experiments and debug training/inference stability.
- **Week 3:** Pixel-space vs. latent-space comparison (or scoped ablation if latent-space implementation is incomplete).
- **Week 4:** Final experiments, analysis, figures/tables, and report writing in conference-paper format.

Risk and Scope Management

The main technical risk is implementation complexity for adapting drifting models to inverse optimization and for building a stable latent-space pipeline. If needed, we will reduce scope by (1) prioritizing one inverse task, (2) using only MNIST/CIFAR-10, and (3) treating latent-space comparison as an ablation/feasibility study rather than a full benchmark.

References

[1] Mingyang Deng, He Li, Tianhong Li, Yilun Du, and Kaiming He. Generative Modeling via Drifting. arXiv preprint arXiv:2602.04770, 2026. `https://arxiv.org/pdf/2602.04770`.

[2] Boyang Zheng, Nanye Ma, Shengbang Tong, and Saining Xie. Diffusion Transformers with Representation Autoencoders. arXiv preprint arXiv:2510.11690, 2025. `https://arxiv.org/pdf/2510.11690`.

[3] Giannis Daras, Hyungjin Chung, Chieh-Hsin Lai, Yuki Mitsufuji, Jong Chul Ye, Peyman Milanfar, Alexandros G. Dimakis, and Mauricio Delbracio. A Survey on Diffusion Models for Inverse Problems. arXiv preprint arXiv:2410.00083, 2024. `https://arxiv.org/pdf/2410.00083`.