Rectified Diffusion: Straightness Is Not Your Need in Rectified Flow

Fu-Yun Wang[1] Ling Yang[2] Zhaoyang Huang[1] Mengdi Wang[3] Hongsheng Li[1]

[1]CUHK MMLab [2]Peking University [3]Princeton University

ICLR International Conference On Learning Representations

Motivation

- Rectified flow, a widely recognized solution, improves generation speed by straightening the ODE path.
- The proposed "rectification" technique for enhancing the quality of few-step generation is deemed viable solely for linear interpolation forms.
- This paper investigates what is most essential about rectified flow.

Training Efficiency Comparison

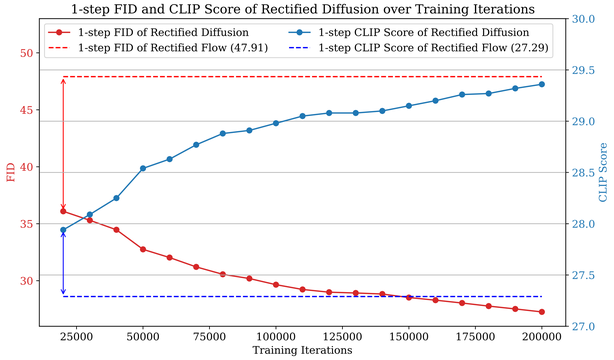

1-step performance of rectified diffusion significantly surpasses the 1-step performance of rectified flow within only 20,000 iterations with batch size 128 (only 8% trained images of rectified flow) and consistently grows with more training iterations.

Effectiveness of Classifier-Free Guidance.

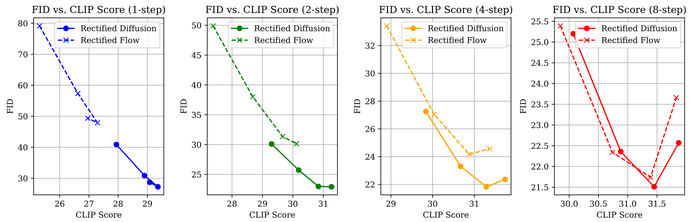

Connection to Diffusion Models and Consistency Models.

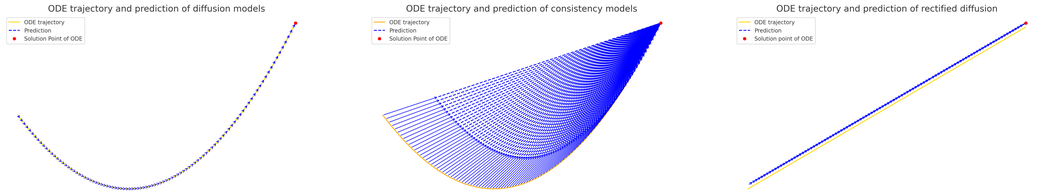

ODE trajectory and prediction comparison of consistency models and rectified diffusion. Since it is difficult to visually tell whether a curved ODE path satisfies first-order approximate property, we apply straight lines for more clear demonstration.

Method Comparison

- Rectified diffusion extending the application form of rectified flow by identifying the straightness is not the essential training target of rectified flow. Overall, rectified diffusion is simpler, stronger, and more general.
- **Simpler**: Rectified diffusion keeps everything of the pretrained diffusion models unchanged, including noise schedulers, prediction types, network architectures, and even training and inference code.
- **Stronger**: We have no gap in prediction type and diffusion form gap between the pretraining and retraining (rectification). This allows for great improvements in training efficiency.
- **General**: We show that rectification is suitable for general diffusion forms, which extends the scope of rectification for general diffusion forms.

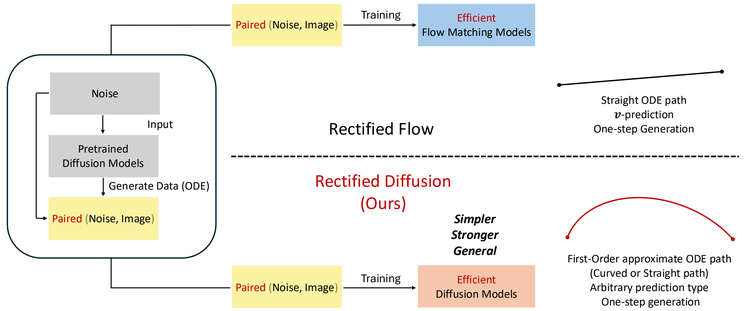

We contend that rectified flow's efficacy arises from leveraging a pretrained diffusion model to obtain noise-sample pairs, subsequently retrained with these pairs. Moreover, we assert that straightness becomes a nonessential training objective when expanding the design space beyond linear interpolating diffusion to broader diffusion forms. Instead, we propose the critical training goal is achieving a first-order approximate ODE path, minimizing high-order discretization errors during inference.

Qualitative Comparison