Adaptive Rectified Flow via Time-Dependent Reweighting

Ji Yu Han, Jung Woo Lee*

Seoul National Univ., *Seoul National Univ.

guhan53@snu.ac.kr, *junglee@snu.ac.kr

시간 의존적 재가중치를 이용한 적응형 Rectified Flow

한지유, 이정우*

서울대학교, *서울대학교

Abstract

Rectified Flow enables efficient generative modeling through deterministic transport between noise and data distributions. However, the standard training procedure applies uniform weighting across all timesteps, implicitly assuming that each phase of the transport process is equally important. In this paper, we propose a simple time-dependent reweighting strategy that emphasizes intermediate timesteps during training, where major structural transformations empirically occur.

Experiments on CIFAR-100 demonstrate substantial improvements in fast sampling performance. Midpoint reweighting achieves 45.8% improvement at five function evaluations and 37.8% at ten evaluations compared to baseline Rectified Flow, while maintaining identical inference cost. Trajectory analysis reveals smoother transport paths with reduced curvature. These results suggest that adaptive temporal emphasis can improve generation efficiency without modifying model architecture or sampling procedures.

Ⅰ. Introduction

Flow-based generative models provide a deterministic alternative to diffusion-based approaches [2, 3], offering fast sampling and stable training. Rectified Flow [1] formulates probability transport between a noise distribution and a data distribution as an ordinary differential equation (ODE) and learns a velocity field that follows straight-line trajectories between paired samples. This formulation simplifies training to a regression problem and enables efficient generation with a small number of integration steps.

Standard Rectified Flow employs a uniformly weighted training objective over the continuous time interval. This design assumes all timesteps contribute equally to learning the transport dynamics. But, the generation process is not uniform over time. Early timesteps primarily involve noise initialization, while late timesteps refine fine-scale details. In contrast, intermediate timesteps often correspond to critical transitions where global structure emerges.

Recent work has explored improving reflow through better sample selection [4], while we investigate whether emphasizing specific temporal regions can improve learned transport trajectories. Our time-dependent reweighting strategy assigns higher importance to mid-trajectory timesteps without changing the model architecture or inference procedure.

Experiments on CIFAR-100 show that this produces smoother trajectories with reduced curvature and significantly improves sample quality in low-NFE regimes.

Ⅱ. Method

Rectified Flow [1] learns a time-dependent velocity field that deterministically transports samples from a noise distribution to a target data distribution. During training, we construct a noisy input sample by linearly interpolating between a randomly drawn noise sample and a data sample using a continuous time variable between zero and one. The velocity network minimizes the mean squared error between the predicted velocity and the straight-line transport direction, with timesteps sampled uniformly over the interval. At inference time, we generate samples by starting from random noise and numerically integrating the learned velocity field using a fixed number of discretization steps, referred to as the number of function evaluations (NFE).

We introduce a time-dependent reweighting strategy that assigns different importance to different temporal regions, modifying only the training objective. The per-sample squared error is multiplied by a scalar weight determined by the corresponding timestep before averaging over the batch. In this work, we focus on a midpoint reweighting strategy that emphasizes intermediate timesteps while down-weighting early and late stages of the transport process. The weighting function consists of a constant baseline term combined with a quadratic component that peaks at the trajectory midpoint.

This differs from standard Rectified Flow by a single multiplicative factor in the loss computation. We make no changes to the model architecture, sampling procedure, or inference-time integration, ensuring identical computational cost during generation.

We evaluate the proposed method on the CIFAR-100 dataset, which has fifty thousand training images and ten thousand test images across one hundred classes. Both baseline and reweighted models use identical U-Net-style architectures with residual blocks and time conditioning. We train for one hundred epochs using the AdamW optimizer with a learning rate of 1e-4 and a cosine annealing schedule. Exponential moving average is applied to model parameters during training.

We assess trajectory quality using a curvature metric that measures variations in the predicted velocity field along the transport path. Sample quality is evaluated under varying NFE settings ranging from five to one hundred steps using reconstruction-based error metrics. Results show that midpoint reweighting reduces trajectory curvature by 5.3% compared to baseline and yields substantial improvements in low-NFE regimes: 45.8% improvement at five steps and 37.8% at ten steps. As NFE increases, the performance gap narrows and both methods converge to similar quality, indicating that the proposed method primarily improves trajectory discretization rather than model capacity.

Ⅲ. Conclusion

We presented a simple time-dependent reweighting strategy for Rectified Flow that emphasizes intermediate timesteps during training. The proposed method requires only a minimal modification to the loss function, introduces no inference-time overhead, and preserves the original model architecture and sampling procedure.

Experiments on CIFAR-100 show that midpoint reweighting significantly improves sample quality in low-NFE regimes by producing smoother transport trajectories with reduced curvature. Not all timesteps are equally important in generative flow training-adaptive temporal emphasis can lead to more efficient generation.

Future work includes extending this approach to higher-resolution datasets, exploring alternative or learnable weighting functions, and developing theoretical analyses connecting temporal reweighting to ODE discretization error.

ACKNOWLEDGMENT

This work is in part supported by the National Research Foundation of Korea (NRF, RS-2024-00451435(20%), RS-2024-00413957(20%)), Institute of Information & communications Technology Planning & Evaluation (IITP, RS-2025-02305453(15%), RS-2025-02273157(15%), RS-2025-25442149(15%) RS-2021-II211343(15%)) grant funded by the Ministry of Science and ICT (MSIT), Institute of New Media and Communications(INMAC), and the BK21 FOUR program of the Education, Artificial Intelligence Graduate School Program (Seoul National University), and Research Program for Future ICT Pioneers, Seoul National University in 2026.

REFERENCES

[1] X. Liu, C. Gong, and Q. Liu, “Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow,” International Conference on Learning Representations, 2023.

[2] J. Ho, A. Jain, and P. Abbeel, “Denoising Diffusion Probabilistic Models,” Neural Information Processing Systems, 2020.

[3] Y. Song et al., “Score-Based Generative Modeling through Stochastic Differential Equations,” International Conference on Learning Representations, 2021.

[4] S. Kim, M. Kwon, J. Jeong, and Y. Uh, “Balanced Conic Rectified Flow”, Neural Information Processing Systems, 2025.