

References |

-  Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden, **Stochastic interpolants: A unifying framework for flows and diffusions**, arXiv preprint arXiv:2303.08797 (2023).
-  Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg, **Structured denoising diffusion models in discrete state-spaces**, Advances in Neural Information Processing Systems 34 (2021), 17981–17993.
-  Avishek Joey Bose, Tara Akhound-Sadegh, Guillaume Huguet, Kilian Fatras, Jarrid Rector-Brooks, Cheng-Hao Liu, Andrei Cristian Nica, Maksym Korablyov, Michael Bronstein, and Alexander Tong, **Se (3)-stochastic flow matching for protein backbone generation**, arXiv preprint arXiv:2310.02391 (2023).
-  Nicholas M Boffi, Michael S Albergo, and Eric Vanden-Eijnden, **How to build a consistency model: Learning flow maps via self-distillation**, arXiv preprint arXiv:2505.18825 (2025).
-  Nicholas Matthew Boffi, Michael Samuel Albergo, and Eric Vanden-Eijnden, **Flow map matching with stochastic interpolants: A mathematical framework for consistency models**, Transactions on Machine Learning Research (2025).
-  Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Soumyadip Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein, **Universal guidance for diffusion models**, Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2023, pp. 843–852.
-  Andrew Campbell, Joe Benton, Valentin De Bortoli, Thomas Rainforth, George Deligiannidis, and Arnaud Doucet, **A continuous time framework for discrete denoising models**, Advances in Neural Information Processing Systems 35 (2022), 28266–28279.
-  Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye, **Diffusion posterior sampling for general noisy inverse problems**, arXiv preprint arXiv:2209.14687 (2022).

References II

-  Itay Chachy, Guy Yariv, and Sagie Benaim, **Rewardsds: Aligning score distillation via reward-weighted sampling**, arXiv preprint arXiv:2503.09601 (2025).
-  Sander Dieleman, **Noise schedules considered harmful**, 2024.
-  Prafulla Dhariwal and Alexander Nichol, **Diffusion models beat GANs on image synthesis**, Advances in Neural Information Processing Systems 34 (2021).
-  Kevin Frans, Danijar Hafner, Sergey Levine, and Pieter Abbeel, **One step diffusion via shortcut models**, arXiv preprint arXiv:2410.12557 (2024).
-  Jiajun Fan, Shuaike Shen, Chaoran Cheng, Yuxin Chen, Chumeng Liang, and Ge Liu, **Online reward-weighted fine-tuning of flow matching with wasserstein regularization**, The Thirteenth International Conference on Learning Representations, 2025.
-  Zhengyang Geng, Mingyang Deng, Xingjian Bai, J Zico Kolter, and Kaiming He, **Mean flows for one-step generative modeling**, arXiv preprint arXiv:2505.13447 (2025).
-  Ruiqi Gao, Emiel Hoogeboom, Jonathan Heek, Valentin De Bortoli, Kevin P. Murphy, and Tim Salimans, **Diffusion meets flow matching: Two sides of the same coin**, 2024.
-  Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and Yaron Lipman, **Discrete flow matching**, Advances in Neural Information Processing Systems 37 (2024), 133345–133385.
-  Jonathan Ho, Ajay Jain, and Pieter Abbeel, **Denoising diffusion probabilistic models**, Advances in Neural Information Processing Systems 33 (2020), 6840–6851.
-  Xixi Hu, Runlong Liao, Keyang Xu, Bo Liu, Yeqing Li, Eugene Ie, Hongliang Fei, and Qiang Liu, **Improving rectified flow with boundary conditions**, arXiv preprint arXiv:2506.15864 (2025).

References III

-  Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling, **Argmax flows and multinomial diffusion: Learning categorical distributions**, Advances in neural information processing systems 34 (2021), 12454–12465.
-  Xixi Hu, Keyang Xu, Bo Liu, Qiang Liu, and Hongliang Fei, **Amo sampler: Enhancing text rendering with overshooting**, Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR), June 2025, pp. 13157–13166.
-  Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine, **Elucidating the design space of diffusion-based generative models**, Advances in neural information processing systems 35 (2022), 26565–26577.
-  Tero Karras, Miika Aittala, Tuomas Kynkänniemi, Jaakko Lehtinen, Timo Aila, and Samuli Laine, **Guiding a diffusion model with a bad version of itself**, Advances in Neural Information Processing Systems 37 (2024), 52996–53021.
-  Diederik P. Kingma and Ruiqi Gao, **Understanding diffusion objectives as the elbo with simple data augmentation**, 2023.
-  Rafal Karczewski, Markus Heinonen, and Vikas Garg, **Diffusion models as cartoonists: The curious case of high density regions**, Proceedings of the International Conference on Learning Representations (ICLR), 2025, Published as a conference paper at ICLR 2025.
-  Dongjun Kim, Chieh-Hsin Lai, Wei-Hsiang Liao, Naoki Murata, Yuhta Takida, Toshimitsu Uesaka, Yutong He, Yuki Mitsufuji, and Stefano Ermon, **Consistency trajectory models: Learning probability flow ode trajectory of diffusion**, arXiv preprint arXiv:2310.02279 (2023).
-  Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le, **Flow matching for generative modeling**, arXiv preprint arXiv:2210.02747 (2022).

References IV

-  Grace Luo, Trevor Darrell, Oliver Wang, Dan B Goldman, and Aleksander Holynski, **Readout guidance: Learning control from diffusion features**, Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 8217–8227.
-  Weijian Luo, Zemin Huang, Zhengyang Geng, J Zico Kolter, and Guo-jun Qi, **One-step diffusion distillation through score implicit matching**, Advances in Neural Information Processing Systems 37 (2024), 115377–115408.
-  Yihong Luo, Tianyang Hu, Weijian Luo, Kenji Kawaguchi, and Jing Tang, **Reward-instruct: A reward-centric approach to fast photo-realistic image generation**, arXiv preprint arXiv:2503.13070 (2025).
-  Weijian Luo, Tianyang Hu, Shifeng Zhang, Jiacheng Sun, Zhenguo Li, and Zhihua Zhang, **Diff-instruct: A universal approach for transferring knowledge from pre-trained diffusion models**, Advances in Neural Information Processing Systems 36 (2023), 76525–76546.
-  Qiang Liu, **Rectified flow: A marginal preserving approach to optimal transport**, arXiv preprint arXiv:2209.14577 (2022).
-  Qiang Liu, **Let us flow together**, arXiv preprint (2024).
-  Cheng Lu and Yang Song, **Simplifying, stabilizing and scaling continuous-time consistency models**, arXiv preprint arXiv:2410.11081 (2024).
-  Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B Hashimoto, **Diffusion-Im improves controllable text generation**, arXiv preprint arXiv:2205.14217 (2022).
-  Weijian Luo, **Diff-instruct++: Training one-step text-to-image generator model to align with human preferences**, arXiv preprint arXiv:2410.18881 (2024).

References V

-  Xingchao Liu, Lemeng Wu, Mao Ye, and Qiang Liu, [Let us build bridges: Understanding and extending diffusion generative models](#), arXiv preprint arXiv:2208.14699 (2022).
-  Wenze Liu and Xiangyu Yue, [Learning to integrate diffusion odes by averaging the derivatives](#), arXiv preprint arXiv:2505.14502 (2025).
-  Xingchao Liu, Xiwen Zhang, Jianzhu Ma, Jian Peng, et al., [Instaflow: One step is enough for high-quality diffusion-based text-to-image generation](#), The Twelfth International Conference on Learning Representations, 2023.
-  Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro, [Exploring generalization in deep learning](#), Advances in neural information processing systems 30 (2017).
-  Alexander Quinn Nichol and Prafulla Dhariwal, [Improved denoising diffusion probabilistic models](#), International Conference on Machine Learning, PMLR, 2021, pp. 8162–8171.
-  Thuan Hoang Nguyen and Anh Tran, [Swiftbrush: One-step text-to-image diffusion model with variational score distillation](#), Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 7807–7816.
-  Matthew Niedoba, Berend Zwartenberg, Kevin Murphy, and Frank Wood, [Towards a mechanistic explanation of diffusion model generalization](#), arXiv preprint arXiv:2411.19339 (2024).
-  Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-Rong Wen, and Chongxuan Li, [Large language diffusion models](#), arXiv preprint arXiv:2502.09992 (2025).
-  Jingyang Ou, Shen Nie, Kaiwen Xue, Fengqi Zhu, Jiacheng Sun, Zhenguo Li, and Chongxuan Li, [Your absorbing discrete diffusion secretly models the conditional distributions of clean data](#), arXiv preprint arXiv:2406.03736 (2024).

References VI

-  Orlando Romero and Mouhacine Benosman, **Finite-time convergence in continuous-time optimization**, International conference on machine learning, PMLR, 2020, pp. 8200–8209.
-  Saptarshi Roy, Vansh Bansal, Purnamrita Sarkar, and Alessandro Rinaldo, **2-rectifications are enough for straight flows: A theoretical insight into wasserstein convergence**.
-  Litu Rout, Yujia Chen, Nataniel Ruiz, Constantine Caramanis, Sanjay Shakkottai, and Wen-Sheng Chu, **Semantic image inversion and editing using rectified stochastic differential equations**, arXiv preprint arXiv:2410.10792 (2024).
-  Allen Z Ren, Justin Lidard, Lars L Ankile, Anthony Simeonov, Pulkit Agrawal, Anirudha Majumdar, Benjamin Burchfiel, Hongkai Dai, and Max Simchowitz, **Diffusion policy policy optimization**, arXiv preprint arXiv:2409.00588 (2024).
-  Subham Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin Chiu, Alexander Rush, and Volodymyr Kuleshov, **Simple and effective masked diffusion language models**, Advances in Neural Information Processing Systems 37 (2024), 130136–130184.
-  Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever, **Consistency models**.
-  Subham Sekhar Sahoo, Justin Deschenaux, Aaron Gokaslan, Guanghan Wang, Justin Chiu, and Volodymyr Kuleshov, **The diffusion duality**, arXiv preprint arXiv:2506.10892 (2025).
-  Amirmojtaba Sabour, Sanja Fidler, and Karsten Kreis, **Align your flow: Scaling continuous-time flow map distillation**, arXiv preprint arXiv:2506.14603 (2025).
-  Tim Salimans and Jonathan Ho, **Progressive distillation for fast sampling of diffusion models**, arXiv preprint arXiv:2202.00512 (2022).

References VII

-  Hannes Stark, Bowen Jing, Chenyu Wang, Gabriele Corso, Bonnie Berger, Regina Barzilay, and Tommi Jaakkola, **Dirichlet flow matching with applications to dna sequence design**, arXiv preprint arXiv:2402.05841 (2024).
-  Jiaming Song, Chenlin Meng, and Stefano Ermon, **Denoising diffusion implicit models**, International Conference on Learning Representations, 2020.
-  Neta Shaul, Juan Perez, Ricky TQ Chen, Ali Thabet, Albert Pumarola, and Yaron Lipman, **Bespoke solvers for generative flow models**, arXiv preprint arXiv:2310.19075 (2023).
-  Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole, **Score-based generative modeling through stochastic differential equations**, International Conference on Learning Representations, 2020.
-  Ravid Shwartz-Ziv and Naftali Tishby, **Opening the black box of deep neural networks via information**, arXiv preprint arXiv:1703.00810 (2017).
-  Utkarsh Utkarsh, Pengfei Cai, Alan Edelman, Rafael Gomez-Bombarelli, and Christopher Vincent Rackauckas, **Physics-constrained flow matching: Sampling generative models with hard constraints**, arXiv preprint arXiv:2506.04171 (2025).
-  Masatoshi Uehara, Yulai Zhao, Tommaso Biancalani, and Sergey Levine, **Understanding reinforcement learning-based fine-tuning of diffusion models: A tutorial and review**, arXiv preprint arXiv:2407.13734 (2024).
-  Masatoshi Uehara, Yulai Zhao, Chenyu Wang, Xiner Li, Aviv Regev, Sergey Levine, and Tommaso Biancalani, **Inference-time alignment in diffusion models with reward-guided generation: Tutorial and review**, arXiv preprint arXiv:2501.09685 (2025).

References VIII

-  Yifei Wang, Weimin Bai, Colin Zhang, Debing Zhang, Weijian Luo, and He Sun, **Uni-instruct: One-step diffusion model through unified diffusion divergence instruction**, arXiv preprint arXiv:2505.20755 (2025).
-  Zhangkai Wu, Xuhui Fan, Hongyu Wu, and Longbing Cao, **Traflow: Trajectory distillation on pre-trained rectified flow**, arXiv preprint arXiv:2502.16972 (2025).
-  Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu, **Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation**, Advances in neural information processing systems 36 (2023), 8406–8441.
-  Chenguang Wang, Xiaoyu Zhang, Kaiyuan Cui, Weichen Zhao, Yongtao Guan, and Tianshu Yu, **Importance weighted score matching for diffusion samplers with enhanced mode coverage**, arXiv preprint arXiv:2505.19431 (2025).
-  Yilun Xu, Weili Nie, and Arash Vahdat, **One-step diffusion models with f -divergence distribution matching**, arXiv preprint arXiv:2502.15681 (2025).
-  Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Fredo Durand, and Bill Freeman, **Improved distribution matching distillation for fast image synthesis**, Advances in neural information processing systems 37 (2024), 47455–47487.
-  Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T Freeman, and Taesung Park, **One-step diffusion with distribution matching distillation**, Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2024, pp. 6613–6623.
-  Longlin Yu, Jiajun Zha, Tong Yang, Tianyu Xie, Xiangyu Zhang, S-H Gary Chan, and Cheng Zhang, **Continuous semi-implicit models**, arXiv preprint arXiv:2506.06778 (2025).

References IX



Mingyuan Zhou, Huangjie Zheng, Zhendong Wang, Mingzhang Yin, and Hai Huang, **Score identity distillation: Exponentially fast distillation of pretrained diffusion models for one-step generation**, Forty-first International Conference on Machine Learning, 2024.