PATH INVARIANCE AND THE ROBUSTNESS OF FLOW MATCHING: BEYOND ARCHITECTURAL AND DATA PERTURBATIONS

Eshant English[1,*] **Taiji Suzuki**[1,2]
[1]University of Tokyo, Tokyo, Japan
[2]RIKEN-AIP, Tokyo, Japan

ABSTRACT

Recent advances in Flow Matching (FM) models have demonstrated remarkable empirical success, often attributed to their ability to learn stable velocity fields that transport probability distributions from noise to data. Recent work has further revealed a striking form of global stability: FM models trained under severe perturbations, such as architectural changes or training on disjoint subsets of data, can still generate visually similar samples when initialised from the same noise realisation. While these observations suggest that FM models learn a robust latent-to-data geometry, the precise origin of this stability remains unclear. In this work, we investigate whether this stability arises from the specific probability path used during training, potentially through a mechanism that "locks" trajectories onto the data manifold. To test this hypothesis, we systematically vary the probability path while fixing the initial latent seed. We compare Linear Optimal Transport (OT), Variance Preserving (VP), Linear VP, and Trigonometric/Cosine paths on the MNIST dataset. Our results show that trajectories initialised from the same noise converge to perceptually similar samples along different paths, ruling out the hypothesis that stability is primarily induced by a particular interpolant. At the same time, we observe clear differences in optimisation behaviour across paths. OT-based paths converge reliably and stably, consistent with recent empirical practice, whereas Linear VP paths exhibit pronounced instability during training. These findings suggest a nuanced conclusion: while the final generative mapping appears largely invariant to the choice of probability path, the path remains crucial from the perspective of optimisation and generalisation. Our results indicate that Flow Matching learns a robust global geometry of the data manifold, while the probability path mainly determines how easily this geometry can be learned by a neural network.

1 INTRODUCTION

Generative modelling Goodfellow et al. (2016); Yang et al. (2025); Zhai et al. (2024) has undergone a significant shift with the emergence of diffusion-based methods (Song et al., 2020) and, more recently, Flow Matching (FM) (Lipman et al., 2022; Albergo & Vanden-Eijnden, 2022; Albergo et al., 2025) frameworks. Unlike traditional diffusion models that rely on stochastic dynamics and entropic transport objectives, Flow Matching learns a deterministic velocity field that defines an Ordinary Differential Equation (ODE) transporting samples from a simple noise distribution to a complex data distribution. This formulation offers conceptual simplicity, competitive sample quality, and improved sampling efficiency.

A growing body of empirical evidence suggests that Flow Matching models exhibit a distinctive form of *global stability* (Briq et al., 2025; Bertrand et al., 2025). Recent studies have demonstrated that FM models remain stable even under severe perturbations, such as aggressive data pruning, reduced model capacity, or training on entirely disjoint subsets of the same dataset (Briq et al., 2025). A particularly striking observation is that trajectories initialised from the same random seed

[*]contact email: eshantenglish@g.ecc.u-tokyo.ac.jp

tend to evolve into visually similar samples, regardless of the specific training subset or architectural configuration. This behaviour suggests that the latent-to-data mapping learned by FM models is governed by a robust global structure.

However, while stability with respect to data and architecture is increasingly well documented, the role of the probability path, the specific interpolation between the noise and data distributions, remains comparatively underexplored. In Flow Matching, the choice of path determines the intermediate distributions, the curvature of trajectories, and the temporal structure of the velocity field. If the observed stability is a truly global property of the learned manifold, one would expect the final generated samples to remain invariant not only to perturbations in data or architecture, but also to the functional form of the probability path itself.

In this work, we explicitly test this hypothesis. We ask whether the stability of Flow Matching arises from a particular interpolant that implicitly constrains trajectories to lock onto the data manifold. To this end, we compare several commonly used probability paths, Linear Optimal Transport (OT), Variance Preserving (VP), Linear VP, and Trigonometric/Cosine paths, while fixing the initial latent samples during inference. By analysing the resulting generations, we aim to assess the degree of path invariance in FM models and to understand the role of the probability path in both stability and optimisation.

2 Preliminaries: Flow Matching Framework

Flow Matching (FM) defines a continuous-time transformation between a noise distribution p_0 and a data distribution p_1. Given a noise sample $x_0 \sim \mathcal{N}(0, I)$ and a data sample $x_1 \sim p_1$, we construct a time-dependent path x_t using scalar coefficients α_t and σ_t for $t \in [0, 1]$.

2.1 Interpolant Construction

The interpolant is defined as a linear combination of the endpoints:

$$x_t = \alpha_t x_1 + \sigma_t x_0. \tag{1}$$

The dynamics of the trajectory are governed by the choice of interpolant path, which specifies α_t and σ_t as well as their time derivatives $\dot{\alpha}_t$ and $\dot{\sigma}_t$.

2.2 Probability Paths and interpolant paths

For any sample pair (x_0, x_1), the conditional velocity field induced by the interpolant is given by

$$\dot{x}_t = \dot{\alpha}_t x_1 + \dot{\sigma}_t x_0. \tag{2}$$

Different choices of α_t and σ_t define different probability paths and induce distinct velocity field geometries. In this work, we consider the following interpolant paths (Lipman et al., 2024; Tong et al., 2023):

- **Conditional OT (CondOT)**: A linear path with $\alpha_t = t$ and $\sigma_t = 1 - t$, corresponding to a constant velocity $x_1 - x_0$.
- **Variance Preserving (VP)**: A curved path derived from diffusion-based schedules, with $\alpha_t = e^{-0.5T(t)}$ and $\sigma_t = \sqrt{1 - e^{-T(t)}}$.
- **Linear VP**: A simplified variance-preserving path defined by $\alpha_t = t$ and $\sigma_t = \sqrt{1 - t^2}$.
- **Trignometric/Cosine**: A trigonometric path with $\alpha_t = \sin(\frac{\pi}{2}t)$ and $\sigma_t = \cos(\frac{\pi}{2}t)$.

2.3 Flow Matching Objective and Inference

The velocity model $v_\theta(x, t)$ is trained by minimising a regression loss that matches the target velocity:

$$\mathcal{L}_{FM}(\theta) = \mathbb{E}_{t,x_0,x_1}\left[\|v_\theta(x_t, t) - (\dot{\alpha}_t x_1 + \dot{\sigma}_t x_0)\|^2\right]. \tag{3}$$

At inference time, samples are generated by integrating the probability flow ODE,

$$\frac{dx}{dt} = v_\theta(x, t), \tag{4}$$

starting from $x(0) = x_0$. Prior work has shown that FM models are stable under perturbations to training data and architecture. We extend these observations by examining whether trajectories initialised from the same x_0 remain invariant across fundamentally different probability paths.

3 Experiments and Results

We evaluate path invariance by training a fixed U-Net architecture (Ronneberger et al., 2015) on the MNIST dataset for 100 epochs using each path.

3.1 Methodology

To isolate the effect of the probability path, we fix the initial latent samples x_0 across all models during inference. This enables a direct comparison of how the same point in the noise distribution is transported to the data manifold under different velocity field geometries.

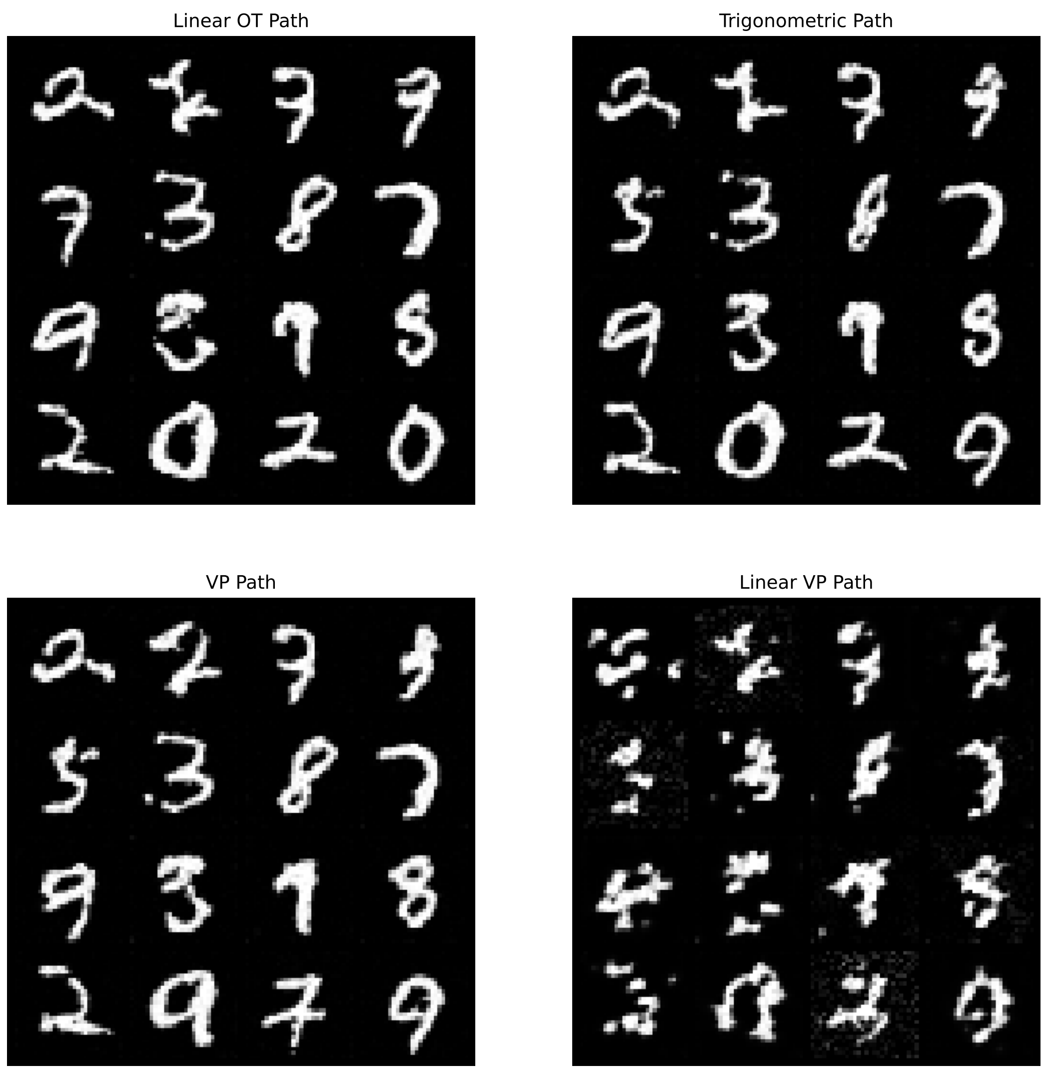

Figure 1: Comparison of generated MNIST samples across four probability paths. Trajectories initialised from the same noise sample x_0 converge to visually similar outputs x_1, demonstrating path invariance. While the Linear OT path, the Trigonometric/Cosine path, and the VP Path show high similarity, the Linear VP path showed noisy generation; however, closer inspection shows the generated shapes are similar to the ones from the other three paths.

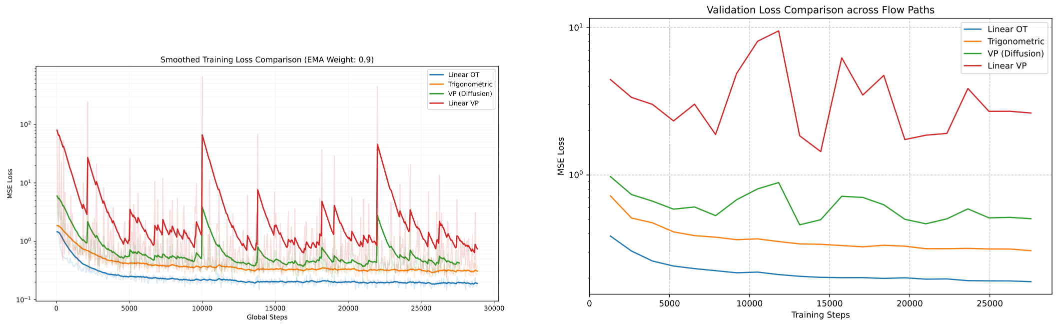

(a) Smoothed training loss (EMA, $\beta = 0.9$).

(b) Validation loss per epoch.

Figure 2: Training and validation performance across different probability paths. The LinearOT path shows better optimisation for the given architecture, whereas the Linear VP path shows a very wiggly behaviour. VP (Diffusion) path is slightly unstable, whereas the Trigonometric/Cosine interpolation is stable but with a higher loss than the OT path.

3.2 Observations

Despite each interpolant path inducing a distinct intermediate probability path and a different temporal structure for the velocity field, the final generated digits are perceptually nearly identical across most paths (see Figure 1). This indicates that Flow Matching converges to a consistent latent-to-data mapping that is robust to the choice of interpolant.

Importantly, this invariance holds even though the instantaneous velocities $\dot{x}_t$ differ substantially across paths. This suggests that the model adapts locally to satisfy the imposed dynamics while preserving a shared global structure. At the same time, we observe pronounced differences in training behaviour: OT-based paths converge reliably and stably, whereas Linear VP paths exhibit significant instability, leading to degraded optimisation performance. These observations are consistent with recent empirical practice favouring OT-style interpolants.

4 Discussion

Our results provide new insight into the origin of stability in Flow Matching models. Prior work (Briq et al., 2025) demonstrated that FM models remain stable under architectural perturbations and training on disjoint subsets of data. In those settings, however, the probability path and intermediate distributions were held fixed, leaving open the possibility that stability arose from a specific transport mechanism.

In contrast, our experiments apply a substantially stronger stress test by explicitly varying the probability path itself. By comparing CondOT, VP, Linear VP, and Trignometric/Cosine interpolants (Albergo et al., 2025; Lipman et al., 2024), we intentionally alter the intermediate distributions p_t and the velocity field geometry throughout the entire trajectory. Despite these changes, the final generated samples remain largely invariant. This allows us to rule out the hypothesis that stability primarily originates from a particular interpolant that enforces manifold locking.

At the same time, our findings highlight that probability paths remain critically important from the perspective of optimisation and generalisation. While the endpoint mapping appears robust, the ease with which this mapping is learned varies substantially across paths. OT-based interpolants consistently exhibit stable convergence, whereas Linear VP paths introduce significant instability. This suggests that certain paths, or potentially distributions over paths, may be inherently easier for neural networks to optimise, even if they do not alter the final latent-to-data geometry.

Overall, our results point to a nuanced understanding of Flow Matching stability. The global geometry of the learned flow appears to be an intrinsic property of the data manifold and the FM objective, largely independent of the chosen probability path. However, the path plays a crucial role in shaping training dynamics, stability, and generalisation, making it an important object of study beyond endpoint correctness alone.

REFERENCES

Michael Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic Interpolants: A Unifying Framework for Flows and Diffusions. *Journal of Machine Learning Research*, 26(209):1–80, 2025. ISSN 1533-7928. URL http://jmlr.org/papers/v26/23-1605.html.

Michael Samuel Albergo and Eric Vanden-Eijnden. Building Normalizing Flows with Stochastic Interpolants. September 2022. URL https://openreview.net/forum?id=li7qeBbCR1t.

Quentin Bertrand, Anne Gagneux, Mathurin Massias, and Rémi Emonet. On the Closed-Form of Flow Matching: Generalization Does Not Arise from Target Stochasticity, December 2025. URL http://arxiv.org/abs/2506.03719. arXiv:2506.03719 [cs].

Rania Briq, Michael Kamp, Ohad Fried, Sarel Cohen, and Stefan Kesselheim. The Amazing Stability of Flow Matching. November 2025. URL https://openreview.net/forum?id=K6ptzIhoKy.

Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT Press, 2016.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow Matching for Generative Modeling. September 2022. URL https://openreview.net/forum?id=PqvMRDCJT9t.

Yaron Lipman, Marton Havasi, Peter Holderrieth, Neta Shaul, Matt Le, Brian Karrer, Ricky T. Q. Chen, David Lopez-Paz, Heli Ben-Hamu, and Itai Gat. Flow Matching Guide and Code, December 2024. URL http://arxiv.org/abs/2412.06264. arXiv:2412.06264 [cs].

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation, May 2015. URL http://arxiv.org/abs/1505.04597. arXiv:1505.04597 [cs].

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations. October 2020. URL https://openreview.net/forum?id=PxTIG12RRHS&utm_campaign=NLP%20News&utm_medium=email&utm_source=Revue%20newsletter.

Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*, November 2023. ISSN 2835-8856. URL https://openreview.net/forum?id=CD9Snc73AW.

Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion Models: A Comprehensive Survey of Methods and Applications, September 2025. URL http://arxiv.org/abs/2209.00796. arXiv:2209.00796 [cs].

Shuangfei Zhai, Ruixiang Zhang, Preetum Nakkiran, David Berthelot, Jiatao Gu, Huangjie Zheng, Tianrong Chen, Miguel Angel Bautista, Navdeep Jaitly, and Josh Susskind. Normalizing Flows are Capable Generative Models, December 2024. URL https://ui.adsabs.harvard.edu/abs/2024arXiv241206329Z. ADS Bibcode: 2024arXiv241206329Z.