

DELAY FLOW MATCHING

Bolin Zhao^{1,2*}, Xiaoyu Zhang^{2,3*}, Yuting Dong^{1,2}, Xin Lu^{3†}, Wei Lin^{1,2,4,5†}, Qunxi Zhu^{2,4,5†}

¹School of Mathematical Sciences, SCMS, SCAM, and CCSB, Fudan University, China

²Research Institute of Intelligent Complex Systems, Fudan University, China

³College of Systems Engineering, National University of Defense Technology, China

⁴State Key Laboratory of Medical Neurobiology and MOE Frontiers Center for Brain Science, Institutes of Brain Science, Fudan University, China

⁵Shanghai Artificial Intelligence Laboratory, China.

{blzhao21, ytdong22}@m.fudan.edu.cn {wlin, qxzhu}@fudan.edu.cn
{xyzhangphd, xin.lu.lab}@outlook.com

ABSTRACT

Flow matching (FM) based on Ordinary Differential Equations (ODEs) has achieved significant success in generative tasks. However, it faces several inherent limitations, including an inability to model trajectory intersections, capture delay dynamics, and handle transfer between heterogeneous distributions. These limitations often result in a significant mismatch between the modeled transfer process and real-world phenomena, particularly when key coupling or inherent structural information between distributions must be preserved. To address these issues, we propose Delay Flow Matching (DFM), a new FM framework based on Delay Differential Equations (DDEs). Theoretically, we show that DFM possesses *universal approximation capability* for continuous transfer maps. By incorporating delay terms into the vector field, DFM enables trajectory intersections and better captures delay dynamics. Moreover, by designing appropriate initial functions, DFM ensures accurate transfer between heterogeneous distributions. Consequently, our framework preserves essential coupling relationships and achieves more flexible distribution transfer strategies. We validate DFM’s effectiveness across synthetic datasets, single-cell data, and image-generation tasks.

1 INTRODUCTION

Generative modeling is a key and rapidly advancing field in machine learning, focusing on learning transformations between different distributions. It underpins a wide range of applications across many tasks in diverse fields, including image generation (Nichol et al., 2021), molecule design (Sanchez-Lengeling & Aspuru-Guzik, 2018), and single-cell trajectory inference (Saelens et al., 2019). Traditional approaches, such as Variational Autoencoders (Kingma, 2013; Rezende et al., 2014), Generative Adversarial Networks (Goodfellow et al., 2014), and Normalizing Flows (Dinh et al., 2014; 2022; Papamakarios et al., 2021), have achieved notable success, yet they still face challenges like instability and computational inefficiency.

Recently, diffusion models (Sohl-Dickstein et al., 2015; Song & Ermon, 2019; Ho et al., 2020; Song et al., 2020a; Yang et al., 2023) have garnered considerable attention for their ability to model complex data distributions via a diffusion process. These models incrementally add noise to data in a forward process and then learn to reverse this process in order to regenerate data samples, which can be understood as learning stochastic dynamics that interpolate between the prior distribution and the target data distribution (Song et al., 2020b; 2021). Additionally, Flow Matching (FM), a simulation-free method for training continuous normalizing flows (Chen et al., 2018), is proposed to model the transformation between distributions as the flow maps of Neural ODEs (Chen et al., 2018; Grathwohl et al., 2018). FM achieves efficient training by directly regressing on an explicitly

*Equal contribution.

†Corresponding authors.

constructed conditional vector field. Concurrently, the stochastic interpolant (Albergo & Vanden-Eijnden, 2022) and rectified flow (Liu et al., 2023) are introduced, both of which employ flow maps for matching distributions, though from distinct conceptual frameworks.

As discussed above, most existing continuous generative models for distribution transfer rely on ODEs or SDEs. However, ODE-based models have limited representational capacity, restricting their ability to capture a broad range of distribution transfer strategies (Dupont et al., 2019). In contrast, Delay Differential Equations (DDEs) have been widely adopted to model various real-world systems, including neural dynamics (Campbell, 2007), electro-optical systems (Chembo Kouomou et al., 2005), population dynamics (Lotka, 1925), and many biological network motifs (Glass et al., 2021), where delayed feedback mechanisms naturally give rise to DDE-based formulations. To address the limitations of ODEs in dynamic modeling, Neural DDEs (Zhu et al., 2020) and their variants (Ji & Orosz, 2024; Zhu et al., 2022) are proposed, which explicitly incorporate the effects of historical states and thus exhibit a significantly enhanced representational capacity. Despite these advancements, however, there is currently no approach that utilizes the probability flows of DDEs for distribution transport.

Contributions. We present Delay Flow Matching (DFM), a generalized framework that enables more flexible and precise distribution transport strategies using DDEs. The key contributions of this study are summarized as follows:

1. **Development of DFM:** We introduce DFM, a novel generative model framework based on Neural DDEs, which overcomes the inherent limitations of ODE-based models by incorporating delay terms in the vector field and designing appropriate initial functions. DFM can model a broader range of transport strategies, including those with trajectory intersections, and achieves more precise transport between heterogeneous distributions. It also naturally adapts to the probability flow generated by delay dynamical systems.
2. **Theoretical insights:** We rigorously prove that DFM can universally approximate any continuous transport map between source and target distributions. In contrast, ODE-based models cannot represent certain simple transport maps, such as those involving trajectory intersections, and cannot achieve exact transport between heterogeneous distributions.
3. **Integration with advanced techniques:** DFM can be seamlessly integrated with existing methods, such as keypoint-guided optimal transport, enabling more effective alignment with known coupling information during transport.
4. **Empirical validation:** We validate DFM’s effectiveness on both synthetic and real-world datasets. DFM accurately recovers underlying delay dynamics from snapshot data, significantly outperforms existing methods in single-cell trajectory inference, and surpasses ODE-based FM in image generation tasks.

2 PRELIMINARIES

2.1 DDES AND PROBABILITY FLOWS

DDEs are widely used to model systems in which the time evolution depends not only on the current state but also on past states, such as those governed by delayed feedback. For time-dependent DDEs with a single delay term, the general formulation is given by:

$$\begin{aligned} \frac{d\mathbf{x}(t)}{dt} &= \mathbf{u}[t, \mathbf{x}(t), \mathbf{x}(t - \tau)], t \in [0, T], \\ \mathbf{x}(h) &= \psi(h), h \in [-\tau, 0], \end{aligned} \quad (1)$$

where $\mathbf{u}[t, \mathbf{x}(t), \mathbf{x}(t - \tau)] : [0, T] \times \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a smooth vector field which is abbreviated as $\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau)$ in the following, $\psi(h)$ represents the continuous initial function. We further denote $\psi(h; \mathbf{x}_0)$ as the initial function that takes the value \mathbf{x}_0 at time $t = 0$, i.e. $\psi(0; \mathbf{x}_0) = \mathbf{x}_0$.




Figure 1: Comparison of ODE and DDE coupling. (a) A Gaussian distribution is mapped to itself via KPG-OT, where the red lines and arrows denote the prescribed source-target keypoint pairs, enforcing $x \rightarrow -x$. The resulting trajectory intersections cannot be realized by ODE flows, whereas the DDE $\dot{x} = -2x(t - 1)$ with constant initial functions achieves exact coupling. (b) The transport of $\mathcal{U}(-2, 2)$ to $\frac{1}{2}\mathcal{U}(-3, -1) + \frac{1}{2}\mathcal{U}(1, 3)$ violates connectivity preservation of ODE flows, forcing mass near 0 into $(-1, 1)$. DDEs achieve precise transport by assigning distinct initial functions.

For a given initial distribution $\mathbf{x}_0 \sim p_0(\mathbf{x}_0)$ and the corresponding initial functions $\psi(t; \mathbf{x}_0)$, the above DDE (1) induces the associated probability flows $p(\mathbf{x}, t|\psi) : \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}^+$, satisfying the delay Fokker-Planck equation (refer to Appendix B) (Guillouzic et al., 1999; Frank, 2005):

$$\frac{\partial p(\mathbf{x}, t|\psi)}{\partial t} = -\nabla \cdot \left\{ \int d\mathbf{x}_\tau [\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau) \cdot p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi)] \right\}, \quad (2)$$

where $p(\mathbf{x}, t|\psi)$ is the probability density given the initial function, $p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi)$ is the joint probability density of being at \mathbf{x} at time t and at \mathbf{x}_τ at time $t - \tau$. The specific relationships and distinctions between the probability flows of ODEs and DDEs are elaborated in Appendix A.

2.2 OPTIMAL TRANSPORT

Optimal transport (OT) provides a mathematical framework for transforming one probability distribution to another in the most cost-efficient manner (Villani et al., 2009; Santambrogio, 2015; Chen, 2016; Peyré et al., 2019). It seeks the optimal strategy for reallocating mass while minimizing the transportation cost. Formally, given two separable metric spaces \mathcal{X} and \mathcal{Y} with probability measures μ on \mathcal{X} and ν on \mathcal{Y} , the objective is to determine a transport plan π^* that minimizes the total transport cost, as defined in the Kantorovich problem (Kantorovich, 1942):

$$C(\mu, \nu) = \inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathcal{X} \times \mathcal{Y}} c(\mathbf{x}, \mathbf{y}) d\pi(\mathbf{x}, \mathbf{y}), \quad \Pi(\mu, \nu) = \left\{ \pi : P_{\#}^{\mathbf{x}}(\pi) = \mu, P_{\#}^{\mathbf{y}}(\pi) = \nu \right\}, \quad (3)$$

where $c(\mathbf{x}, \mathbf{y})$ denotes the transportation cost of moving a unit mass from \mathbf{x} to \mathbf{y} , $P_{\#}^{\mathbf{x}}(\pi)$ and $P_{\#}^{\mathbf{y}}(\pi)$ denote the marginal distributions of π with respect to \mathbf{x} and \mathbf{y} . In most cases, the transport cost between two points is defined as the squared Euclidean distance. The minimum total transport cost in this scenario corresponds to the squared 2-Wasserstein distance between the two distributions:

$$W_2(\mu, \nu)^2 = \inf_{\pi \in \Pi(\mu, \nu)} \int \|\mathbf{x} - \mathbf{y}\|^2 d\pi(\mathbf{x}, \mathbf{y}). \quad (4)$$

2.3 KEYPOINT-GUIDED OPTIMAL TRANSPORT

Traditional OT, which focuses solely on minimizing transport costs, often neglects other crucial constraints, resulting in suboptimal matching strategies. Thus, it is important to use a small set of well-matched source-target keypoint pairs, $\mathcal{K} = \{(\mathbf{x}_k, \mathbf{y}_k)\}_{k=1}^K$, to semi-supervise the transport strategy, especially when inherent structures and features of data need to be preserved.

To ensure correct transport with these keypoints, Keypoint-guided Optimal Transport (KPG-OT) is proposed (Gu et al., 2022; 2023), which guarantees that the relationships between each data point and the keypoints are preserved during transport. Formally, the objective can be expressed as:

$$\inf_{\tilde{\pi} \in \tilde{\Pi}(\mu, \nu)} \int_{\mathcal{X} \times \mathcal{Y}} g(\mathbf{x}, \mathbf{y}) d(w \circ \tilde{\pi})(\mathbf{x}, \mathbf{y}), \quad \tilde{\Pi}(\mu, \nu) = \left\{ \tilde{\pi} : P_{\#}^{\mathbf{x}}(w \circ \tilde{\pi}) = \mu, P_{\#}^{\mathbf{y}}(w \circ \tilde{\pi}) = \nu \right\}, \quad (5)$$

where $w \circ \tilde{\pi} = w(\mathbf{x}, \mathbf{y})\tilde{\pi}(\mathbf{x}, \mathbf{y})$ is the keypoint-masked transport plan, w denotes the mask function. For a pair of keypoints $(\mathbf{x}_k, \mathbf{y}_k) \in \mathcal{K}$, the mask function is defined as: $w(\mathbf{x}_k, \mathbf{y}_k) = 1$, $w(\mathbf{x}_k, \mathbf{y}) = 0$ for $\mathbf{y} \neq \mathbf{y}_k$, $w(\mathbf{x}, \mathbf{y}_k) = 0$ for $\mathbf{x} \neq \mathbf{x}_k$, and $w(\mathbf{x}, \mathbf{y}) = 1$ if neither \mathbf{x} nor \mathbf{y} matches any keypoint. Obviously, the mask function defined above exactly preserves the matching of the keypoints in \mathcal{K} .

The cost function in Eq. (5) is defined as $g(\mathbf{x}, \mathbf{y}) = d[R^s(\mathbf{x}), R^t(\mathbf{y})]$, where d denotes the Jensen-Shannon divergence. $R^s(\mathbf{x}) \in (0, 1)^K$ (resp. $R^t(\mathbf{y}) \in (0, 1)^K$) captures the relationship between \mathbf{x} (resp. \mathbf{y}) and all source points (resp. target points) in \mathcal{K} , with its k -th dimension representing the relation score between \mathbf{x} (resp. \mathbf{y}) and the k -th source keypoint \mathbf{x}_k (resp. target keypoint \mathbf{y}_k):

$$R_k^s(\mathbf{x}) = \frac{e^{-c(\mathbf{x}, \mathbf{x}_k)/\tau}}{\sum_{i=1}^K e^{-c(\mathbf{x}, \mathbf{x}_i)/\tau}}, \quad R_k^t(\mathbf{y}) = \frac{e^{-c(\mathbf{y}, \mathbf{y}_k)/\tau}}{\sum_{i=1}^K e^{-c(\mathbf{y}, \mathbf{y}_i)/\tau}}, \quad (6)$$

where τ is the temperature hyperparameter. Note that g quantifies the similarity distance between the two relationship vectors outlined above. Therefore, minimizing the cost function in Eq. (5) effectively promotes the preservation of the relationship between each point and the keypoints.

2.4 ODE-BASED FLOW MATCHING

Flow Matching (FM) (Lipman et al., 2022) is a simulation-free training method for CNF, which trains the parameterized vector field $\mathbf{v}(t, \mathbf{x}; \theta)$ by regression on the target vector field $\mathbf{u}(t, \mathbf{x})$ that generates the probability paths $p(\mathbf{x}, t)$ with marginals $p_0 = q_0$ and $p_1 = q_1$. The training objective is

$$\mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}_{t, p(\mathbf{x}, t)} \|\mathbf{v}(t, \mathbf{x}; \theta) - \mathbf{u}(t, \mathbf{x})\|^2, \quad (7)$$

where $t \sim \mathcal{U}(0, 1)$ follows the uniform distribution.

However, the explicit forms of $\mathbf{u}(t, \mathbf{x})$ and $p(\mathbf{x}, t)$ are intractable to compute. To address this, Conditional FM (CFM) (Tong et al., 2023a; Pooladian et al., 2023) introduces a latent variable \mathbf{z} to construct the target probability path as a mixture of conditional probability paths $p(\mathbf{x}, t) = \mathbb{E}_{q(\mathbf{z})}[p(\mathbf{x}, t|\mathbf{z})]$ and modifies the training objective to:

$$\mathcal{L}_{\text{CFM}}(\theta) = \mathbb{E}_{t, q(\mathbf{z}), p(\mathbf{x}, t|\mathbf{z})} \|\mathbf{v}(t, \mathbf{x}; \theta) - \mathbf{u}(t, \mathbf{x}|\mathbf{z})\|^2, \quad (8)$$

where $\mathbf{u}(t, \mathbf{x}|\mathbf{z})$ denotes the conditional vector field generating the conditional probability path $p(\mathbf{x}, t|\mathbf{z})$. Then, it can be shown that the gradient w.r.t θ of the CFM objective (8) is the same as that of the FM objective (7). Typically, we set the latent condition $\mathbf{z} := (\mathbf{x}_0, \mathbf{x}_1)$, with $q(\mathbf{z})$ being the coupling between distributions $q_0(\mathbf{x}_0)$ and $q_1(\mathbf{x}_1)$, based on a coupling strategy such as OT or KPG-OT. The conditional probability path is then modeled as a Gaussian flow:

$$p(\mathbf{x}, t|\mathbf{z}) = \mathcal{N}(\mathbf{x}|t\mathbf{x}_1 + (1-t)\mathbf{x}_0, \sigma^2), \quad (9)$$

with the corresponding conditional vector field:

$$\mathbf{u}(t, \mathbf{x}|\mathbf{z}) = \mathbf{x}_1 - \mathbf{x}_0. \quad (10)$$

3 LIMITATIONS OF ODE-BASED FLOW MATCHING

Current continuous-time generative models utilize the parameterized vector field of ODEs to facilitate distribution transport. However, the inherent constraints of ODEs limit their ability to model specific transport strategies, leading to significant discrepancies between the modeled and true transport processes. In this section, we theoretically demonstrate the limitations of the ODE-based FM frameworks. Note that the following propositions are stated within the classical FM framework (Lipman et al., 2022; Liu et al., 2023). Recent ODE-based FM extensions enhance representational capacity and mitigate these limitations; we discuss these variants further in Section 5.

Proposition 3.1 (Restriction on trajectory intersections). *Suppose that the target transport strategy, guided by certain key points or specific coupling constraints, inevitably induces trajectory intersections in the augmented phase space (t, \mathbf{x}) during the transport process. If the ODE-based FM is governed by a vector field that is Lipschitz continuous in \mathbf{x} and continuous in t , then the corresponding flow map cannot precisely preserve the transport strategy.*

Remark 3.2. As demonstrated by Liu et al. (2023), when trajectory intersections occur in the transport strategy, FM tends to learn a *rectified flow*, with the targets of the trajectories being “rewired” at the intersection points. As a result, the given transport strategy cannot be accurately preserved. A simple example is shown in Fig. 1 (a).

Proposition 3.3 (Heterogeneity in distributions). *Assume that the source (resp., target) distribution q_0 (resp., q_1) is supported on M (resp., N) disjoint compact sets $\{U_i^0\}_{i=1}^M$ (resp., $\{U_j^1\}_{j=1}^N$), where each $U_i^0 \in \mathbb{R}^d$ (resp., $U_j^1 \in \mathbb{R}^d$) is a path-connected set with non-zero measure and $q_0(\mathbf{x}_0) > 0$ (resp., $q_1(\mathbf{x}_1) > 0$) for all $\mathbf{x}_0 \in \bigcup_{i=1}^M U_i^0$ (resp., $\mathbf{x}_1 \in \bigcup_{j=1}^N U_j^1$). If $M < N$, then for any transport map $T : \bigcup_{i=1}^M U_i^0 \rightarrow \bigcup_{j=1}^N U_j^1$ satisfying $T_{\#}q_0 = q_1$, the following holds: if the flow map induced by a Neural ODE is equal to the transport map T , then the associated vector field cannot be Lipschitz continuous.*

Remark 3.4. The conclusion emphasizes that accurate transport using an ODE-based FM requires regularity assumptions on distributions, which are often violated in real-world scenarios. For example, in single-cell dynamics, cells differentiate from one type into multiple types over time, leading to heterogeneous distributions before and after differentiation. A simple example is shown in Fig. 1 (b).

Moreover, when snapshot data are generated by a delay dynamical system, ODE-based FM fails to recover the true vector field with the delay term, resulting in inaccurate distribution transfer, interpolation, and extrapolation predictions.

4 DELAY FLOW MATCHING

To overcome the limitations of ODE-based FM, we introduce the Delay FM (DFM) framework, a new class of generative models based on Neural DDEs for distribution transport. Since the vector field incorporates delay terms, DFM allows trajectory intersections, enabling more accurate transport for tasks requiring keypoint-guided strategies. Additionally, by designing appropriate initial functions, DFM addresses singularities caused by distributional heterogeneity. A comparison between DDE-based and ODE-based FM under various scenarios is shown in Table 1.

Table 1: Comparison of ODE-based and DDE-based FM. DFM can handle not only KPG-OT coupling and trajectory intersections but also distribution heterogeneity by employing diverse initial functions. Furthermore, DFM can accurately model and recover the dynamics of delay dynamical systems.

Methods	KPG-OT	Intersection	Heterogeneity	Delay Dynamics
FM	✗	✗	✗	✗
CFM	✗	✗	✗	✗
DFM(C)	✓	✓	✗	✓
DFM(D)	✓	✓	✓	✓

4.1 FORMULATION

Consider a target probability flow $p(\mathbf{x}, t)$, generated by a vector field with a single delay term $\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau)$ and initial functions $\psi \sim q^\circ(\psi)$. The vector field and initial functions naturally define a joint probability flow at time t and $t - \tau$, denoted as $p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau)$, which satisfies

$$\int p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau) d\mathbf{x}_\tau = p(\mathbf{x}, t). \quad (11)$$

Note that $p(\mathbf{x}, t)$ and $p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau)$ can be modeled as a mixture of conditional distributions $p(\mathbf{x}, t|\psi)$ and $p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi)$, respectively, as follow:

$$\begin{aligned} p(\mathbf{x}, t) &= \mathbb{E}_{q^\circ(\psi)} p(\mathbf{x}, t|\psi), \\ p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau) &= \mathbb{E}_{q^\circ(\psi)} p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi), \end{aligned} \quad (12)$$

where $p(\mathbf{x}, t|\psi)$ and $p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi)$ represent the probability flow generated by $\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau)$ and the initial function ψ , and they satisfy the delay Fokker-Planck equation (2). Based on the joint probability flow, we aim to learn a parameterized vector field $\mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta)$, which can generate the target probability flow, by minimizing the following regression objective against the target vector field:

$$\mathcal{L}_{\text{DFM}}(\theta) = \mathbb{E}_{t, q^\circ(\psi), p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi)} \|\mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta) - \mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau)\|^2. \quad (13)$$

In most cases, both $p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi)$ and $\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau)$ are computationally intractable. To address this, inspired by CFM, we introduce a latent variable \mathbf{z} and further decompose the target probability path and joint probability path conditioned on the initial function into mixtures of simple probability paths and simple joint probability paths conditioned on both ψ and \mathbf{z} , respectively, as follow:

$$p(\mathbf{x}, t|\psi) = \mathbb{E}_{q(\mathbf{z})} p(\mathbf{x}, t|\mathbf{z}, \psi), \quad p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi) = \mathbb{E}_{q(\mathbf{z})} p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\mathbf{z}, \psi), \quad (14)$$

where $p(\mathbf{x}, t|\mathbf{z}, \psi) = \int p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\mathbf{z}, \psi) d\mathbf{x}_\tau$. We can define the marginal vector field by marginalizing over the conditional vector field $\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau|\mathbf{z})$ as follow:

$$\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau) = \mathbb{E}_{q(\mathbf{z})} \frac{\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau|\mathbf{z}) p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\mathbf{z}, \psi)}{p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi)}, \quad (15)$$

where $\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau|\mathbf{z})$ represents the conditional vector field that generates $p(\mathbf{x}, t|\mathbf{z}, \psi)$ and $p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\mathbf{z}, \psi)$. Under this setup, we can derive the following result.

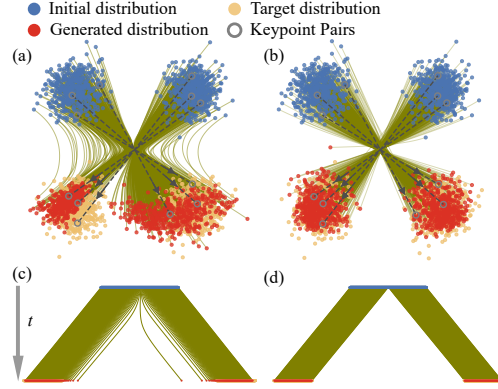


Figure 2: (a, b) show the generation results of CFM (a) and DFM (C) (b) under the KPG-OT coupling on Gaussian mixture data. (c, d) show the generation results of CFM (c) and DFM (D) (d) for heterogeneous source and target distributions.

Proposition 4.1. *The marginal vector field $\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau)$ given in Eq. (15), together with the selected initial function ψ , generates the probability path $p(\mathbf{x}, t|\psi)$ and the joint probability path $p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi)$ in Eq. (14).*

Then, we can train the parameterized vector field by minimizing the following Delay Conditional FM (DCFM) objective:

$$\mathcal{L}_{\text{DCFM}}(\theta) = \mathbb{E}_{t, q(\mathbf{z}), q^\circ(\psi), p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \mathbf{z}, \psi)} \|\mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta) - \mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau | \mathbf{z})\|^2, \quad (16)$$

because the following proposition holds.

Proposition 4.2. *Given that $p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \psi) > 0$ for all $\mathbf{x}, \mathbf{x}_\tau \in \mathbb{R}^d$ and $t \in [0, 1]$, up to a constant independent of θ , $\mathcal{L}_{\text{DCFM}}(\theta)$ and $\mathcal{L}_{\text{DFM}}(\theta)$ are equal, which further implies that $\nabla_\theta \mathcal{L}_{\text{DCFM}}(\theta) = \nabla_\theta \mathcal{L}_{\text{DFM}}(\theta)$.*

4.2 SELECTION OF THE LATENT VARIABLE \mathbf{z}

4.2.1 TRANSPORT BETWEEN TWO MEASURES

As outlined in Section 2.4, ODE-based FM typically selects latent variables as two endpoints $\mathbf{z} := (\mathbf{x}_0, \mathbf{x}_1)$ jointly sampled from the source and target distributions. In contrast, DFM requires the construction of a conditional joint probability density path for $\mathbf{x}(t)$ and $\mathbf{x}(t - \tau)$ based on the latent variable. Hence, we define the latent variable \mathbf{z} as a entire path $\gamma(t; \mathbf{x}_0, \mathbf{x}_1)$ connecting \mathbf{x}_0 and \mathbf{x}_1 .

Formally, we define the probability distribution of the latent variable as $q[\gamma(t; \mathbf{x}_0, \mathbf{x}_1)] := \pi(\mathbf{x}_0, \mathbf{x}_1) \mathcal{P}(\gamma; \mathbf{x}_0, \mathbf{x}_1)$, where $\pi(\mathbf{x}_0, \mathbf{x}_1)$ represents the OT or KPG-OT coupling between the source and target distributions, and $\mathcal{P}(\gamma; \mathbf{x}_0, \mathbf{x}_1)$ denotes the path measure pinned at \mathbf{x}_0 and \mathbf{x}_1 . In practice, we can simply construct the path measure as a Dirac Delta distribution at a given path γ^* connecting \mathbf{x}_0 and \mathbf{x}_1 :

$$\mathcal{P}(\gamma; \mathbf{x}_0, \mathbf{x}_1) = \delta(\gamma - \gamma^*), \quad \gamma_0^* = \mathbf{x}_0, \gamma_1^* = \mathbf{x}_1, \quad (17)$$

where γ^* is constructed using a specific interpolation method. For instance, it can be taken as the linear interpolation $\gamma_t^* = (1 - t)\mathbf{x}_0 + t\mathbf{x}_1$, or alternatively, as a geodesic interpolation on the data manifold, based on appropriate manifold learning techniques.

After sampling a path $\gamma(t; \mathbf{x}_0, \mathbf{x}_1) \sim q(\mathbf{z})$, the conditional joint probability density naturally degenerates into the following form:

$$p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \gamma) = \delta[\mathbf{x} - \gamma(t)] \delta[\mathbf{x}_\tau - \gamma(t - \tau)], \quad (18)$$

which is exactly generated by the conditional vector field:

$$\mathbf{u}[t, \mathbf{x}, \mathbf{x}_\tau | \gamma(t; \mathbf{x}_0, \mathbf{x}_1)] = \frac{\partial \gamma(t; \mathbf{x}_0, \mathbf{x}_1)}{\partial t}. \quad (19)$$

4.2.2 TRANSPORT BETWEEN MULTIPLE PROBABILITY MEASURES OVER TIME

For tasks with multiple target probability distributions over time $\{q_{t_j}\}_{j=0}^J$, where $t_0 = 0, t_J = T$, we define the latent variable $\mathbf{z} := \gamma(t; \{\mathbf{x}_{t_j}\}_{j=0}^J)$ as a trajectory passing through \mathbf{x}_{t_j} at time t_j , sampled from

$$q[\gamma(t; \{\mathbf{x}_{t_j}\}_{j=0}^J)] := \mathcal{P}(\gamma; \{\mathbf{x}_{t_j}\}_{j=0}^J) \prod_{i=0}^{J-1} \pi(\mathbf{x}_{t_i}, \mathbf{x}_{t_{i+1}}), \quad (20)$$

where $\{\mathbf{x}_{t_j}\}_{j=0}^J \sim \prod_{i=0}^{J-1} \pi(\mathbf{x}_{t_i}, \mathbf{x}_{t_{i+1}})$, with $\pi(\mathbf{x}_{t_i}, \mathbf{x}_{t_{i+1}})$ representing the OT or KPG-OT coupling between adjacent probability distributions at time t_i and t_{i+1} , and $\mathcal{P}(\gamma; \{\mathbf{x}_{t_j}\}_{j=0}^J)$ denotes the path measure pinned at $\{\mathbf{x}_{t_j}\}_{j=0}^J$. Similarly, we construct the path measure as:

$$\mathcal{P}(\gamma; \{\mathbf{x}_{t_j}\}_{j=0}^J) = \delta(\gamma - \gamma^*), \quad \gamma_{t_j}^* = \mathbf{x}_{t_j}, \quad (21)$$

where γ^* represents a path passing through \mathbf{x}_{t_j} at time t_j , which can be constructed using the cubic spline (CSpline) interpolation method. With this path as the latent variable, the resulting conditional joint probability path and conditional vector field are identical to those in Eq. (18) and Eq. (19), respectively.

4.3 SELECTION OF THE INITIAL FUNCTION

4.3.1 DFM WITH CONSTANT INITIAL FUNCTIONS

In general, the initial function can be simply chosen as a constant (DFM(C)), i.e. $q^\circ(\psi) = \delta(\psi - \psi^*)$, where $\psi^*(t; \mathbf{x}_0) \equiv \mathbf{x}_0$ for $t \in [-\tau, 0]$ and $\mathbf{x}_0 \sim q_0$. Under this setting, we can theoretically prove that DFM can approximate any continuous transport strategy to arbitrary precision.

Proposition 4.3 (Universal approximating capability of DDE’s flow map). *For any given $\epsilon > 0$ and continuous transport map $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$, which push-forward the source distribution q_0 to the target distribution q_1 , i.e. $F_{\#}q_0 = q_1$, if there exists a neural network $\mathbf{f}(\mathbf{x}; \theta)$ such that $\|\mathbf{f}(\mathbf{x}; \theta) - [F(\mathbf{x}) - \mathbf{x}]\| < \epsilon$, for all $\mathbf{x} \in \mathbb{R}^d$, then we can construct a vector field with a single delay term $\mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta)$, where the corresponding flow map $G(\mathbf{x}; \theta)$, under the constant initial function condition, satisfies $\|G(\mathbf{x}; \theta) - F(\mathbf{x})\| < \epsilon$, for all $\mathbf{x} \in \mathbb{R}^d$.*

This indicates the exceptional representational capacity of DFM, enabling the modeling of a wider range of transport processes than ODE-based FM.

4.3.2 DFM WITH DIVERSE INITIAL FUNCTIONS

In Section 3, we rigorously prove that ODE-based FM cannot effectively handle tasks with heterogeneous source and target distributions. Here, we demonstrate that DFM addresses this by designing diverse initial functions (DFM(D)).

Specifically, we employ clustering methods, such as Gaussian Mixture Model, DBSCAN, to partition the source dataset $\mathbf{X}_0 \sim q_0$ (resp. target dataset $\mathbf{X}_1 \sim q_1$) into M (resp. N) mutually exclusive subsets, denoted as $\mathbf{X}_0^{(1)}, \dots, \mathbf{X}_0^{(M)}$ (resp. $\mathbf{X}_1^{(1)}, \dots, \mathbf{X}_1^{(N)}$). We can assign a normalized mass to each subset $\mathbf{X}_0^{(m)}$ (resp. $\mathbf{X}_1^{(n)}$) as $\rho_0^{(m)} = |\mathbf{X}_0^{(m)}|/|\mathbf{X}_0|$ (resp., $\rho_1^{(n)} = |\mathbf{X}_1^{(n)}|/|\mathbf{X}_1|$). If the endpoints $(\mathbf{x}_0, \mathbf{x}_1)$ of a sampled trajectory γ are drawn from $\mathbf{X}_0^{(m)}$ and $\mathbf{X}_1^{(n)}$, we assign it an initial function ψ_{mn}^* which has a constant time derivative C_{mn} , i.e.:

$$\frac{d\psi_{mn}^*(t; \mathbf{x}_0)}{dt} = C_{mn}, \quad \psi_{mn}^*(0; \mathbf{x}_0) = \mathbf{x}_0, \quad (22)$$

where $t \in [-\tau, 0]$. In this case, $q^\circ(\psi)$ is a discrete distribution which satisfies:

$$\sum_{n=1}^N q^\circ(\psi_{mn}^*) = \rho_0^{(m)}, \quad \sum_{m=1}^M q^\circ(\psi_{mn}^*) = \rho_1^{(n)}. \quad (23)$$

$p(\mathbf{x}_0, 0 | \psi = \psi_{mn}^*) = q_0(\mathbf{x}_0 | \psi = \psi_{mn}^*)$ represents empirical data distribution available as data points in $\mathbf{X}_0^{(m)}$ whose corresponding transport target is in $\mathbf{X}_1^{(n)}$. By coupling the source and target data through OT or KPG-OT, we can obtain the transport target for any initial point. This enables the construction of $q^\circ(\psi)$ as described above, from which we can sample to obtain the corresponding initial function.

In summary, we design distinct initial functions for different subsets of the source and target data, guiding the vector field from different initial subsets to corresponding target subsets, thereby effectively handling distributional heterogeneity.

4.4 GENERATION PROCESS BASED ON NEURAL DDES

After training, we sample data from q_0 and initial functions from $q^\circ(\psi)$, then generate the target data by solving the forward pass of the trained vector field with a single delay term using a piecewise ODE solver, as in Neural DDEs (Zhu et al., 2020).

5 RELATED WORKS

A number of recent methods have been proposed to address the limitations of ODE-based FM. To enable crossing trajectories, Constant Acceleration Flow (Park et al., 2024) models acceleration and conditions it on an initial velocity; Hierarchical Rectified Flow (Zhang et al., 2025) achieves trajectory

crossings by hierarchically coupling multiple ODEs; and Augmented Bridge Matching (De Bortoli et al., 2023) incorporates information about the initial point directly into the vector field. Diversified Flow Matching (Shrestha & Fu, 2025) takes a different perspective by training interpolation strategies that construct non-intersecting target trajectories, thereby indirectly mitigating trajectory-crossing issues. To address the singularities caused by distributional heterogeneity, Switched Flow Matching (Zhu & Lin, 2024) and Variational Rectified Flow Matching (Guo & Schwing, 2025) introduce additional latent variables that allow the model to represent multi-modal transport paths, alleviating the degeneracies that arise under heterogeneous target distributions.

The methods discussed above introduce various extensions to classical FM, enabling the model to accommodate more transport strategies. While these approaches either allow trajectory intersections during the transport process or alleviate singularities arising from distributional heterogeneity, most of them cannot simultaneously resolve both issues. Moreover, they remain confined to the ODE-based generative framework, typically resolving limitations by introducing additional conditioning variables into the vector field or by redesigning the target transport process.

In contrast, our proposed DFM departs from the ODE paradigm and introduces a fundamentally new DDE-based generative framework. By modeling the vector field directly in the original phase space, without introducing auxiliary latent variables or specially engineered transport paths, DFM inherently accommodates both trajectory intersections and distributional heterogeneity within a unified mechanism. Furthermore, DFM naturally captures transport dynamics governed by intrinsic delays, a capability that is entirely out of reach for those ODE-based extensions.

6 EXPERIMENTS

We demonstrate the advantages of DFM over ODE-based FM across various tasks, including reconstructing delay dynamics from snapshots, inferring differentiation trajectories from single-cell RNA sequencing (scRNA-seq) data, and image generation. The experimental details and additional results can be found in Appendix D. A systematic ablation study is shown in Appendix F.

6.1 SIMPLE SYNTHETIC ILLUSTRATIONS

To more clearly illustrate the DFM workflow, we present two simple and concrete running examples, demonstrating the procedure of DFM(C) and DFM(D), respectively.

2-d Gaussian mixture dataset. As shown in Fig. 2 (a,b), we aim to transport a two-component Gaussian mixture to its reflection-symmetric counterpart while respecting 6 keypoint pairs (i.e., the optimal mapping $\mathbf{x} \rightarrow -\mathbf{x}$). Using KPG-OT, we first obtain pointwise couplings and construct the linear interpolation paths $\gamma = t \times \mathbf{x}_0 + (1-t) \times \mathbf{x}_1$, which unavoidably produce trajectory crossings. We select the constant initial function with $\tau = 1$, and train the delay vector field. Notably, this toy problem admits an exact DDE solution: $\dot{\mathbf{x}} = -2\mathbf{x}(t-1) = -2\mathbf{x}_0$. The learned DFM(C) exhibits the same behavior: trajectories select different directions at the crossing by leveraging historical states, yielding accurate keypoint-respecting transport.

1-d example with distributional heterogeneity. As shown in Fig. 2 (c,d), we aim to realize the transport $\mathcal{U}(-1, 1) \rightarrow \frac{1}{2}\mathcal{U}(-3, -2) + \frac{1}{2}\mathcal{U}(2, 3)$. Notably, we can use $\dot{x} = x(t) - x(t-1)$ to achieve precise OT transport by assigning different initial functions: ① $x(t) = x_0 - t, x_0 \in [-1, 0]$ and ② $x(t) = x_0 + t, x_0 \in (0, 1]$. During training, we employed the OT coupling, which maps $[-1, 0]$ to $[-3, -2]$ and $(0, 1]$ to $[2, 3]$, and construct the linear interpolation paths. In applying DFM(D), initial function ① is selected for $x_1 \in [-3, -2]$, while ② is chosen for $x_1 \in [2, 3]$, with $\tau = 1$. The trained DFM(D) achieves accurate transfer between heterogeneous distributions (Fig. 2 (d)).

6.2 SYNTHETIC DATASETS OF DELAY DYNAMICAL SYSTEMS

Biological autoregulation motif. DDEs provide a simplified framework for modeling biological network motifs (Glass et al., 2021). We consider recovering the delay dynamics from snapshot data generated by the autoregulation model: $\dot{x}(t) = \frac{\eta}{1+x^n(t-\tau)}$ with $\tau = 1, n = 2, \eta = 5$, which produces damped oscillations (Fig. 3 (a,b)). Using 1,000 initial values sampled from $\mathcal{U}(0.2, 1.2)$ and constant initial functions, we generate trajectories and collect snapshots at $t=0, 1, \dots, 7$ for training. During




Figure 3: Comparison of predicted trajectories of CFM and DFM trained on snapshots of delay dynamical systems. (a-b) Interpolation and extrapolation (insets) results on the biological autoregulation motif dataset using KP-CFM (a) and KP-DFM(C) (b), both with two keypoints. (c-d) Interpolation results on the spiral DDE using OT-CFM (c) and OT-DFM(C) (d).

training, KPG-OT (KP-) matches minibatches at adjacent time steps using two known keypoints. After training, forward integration from $t=0$ shows that KP-CFM fails to recover the dynamics, while KP-DFM(C) with $\tau=1$ accurately captures and extrapolates the oscillatory behavior (insets in Fig. 3 (a,b)). Results with other delays and error metrics are provided in Appendix D.1 and F.

Spiral DDE. We next consider a 2-d DDE (Zhu et al., 2020): $\dot{x}(t) = \mathbf{A} \tanh[\mathbf{x}(t) + \mathbf{x}(t-\tau)]$ with $\tau=0.5$ and $\mathbf{A} \in \mathbb{R}^{2 \times 2}$, which produces crossing spiral trajectories (Fig. 3 (c,d)). Snapshots are taken every 0.05 in $t \in [0, 0.6]$ and coupled using minibatch-OT (OT-) between adjacent steps. After training, OT-DFM(C) with $\tau=0.5$ successfully reproduces the dynamics, while OT-CFM fails around the crossing region. DFM also generalizes well across a range of τ values (Appendix D.2, F).

6.3 TRAJECTORY INFERENCE OF SINGLE-CELL

We investigate the inference of differentiation trajectories from real scRNA-seq data, where heterogeneity increases as cells transition from a single type to multiple types during development, complicating modeling with ODE-based approaches. To evaluate performance, we train on all time points except one intermediate time, then sample from the initial distribution and perform forward integration to predict distributions at each time point. Predictions are assessed via unsupervised leave-one-out validation (L) against the held-out intermediate distribution and supervised final-time validation (F) against the final distribution. Trajectory accuracy is measured using the 2-Wasserstein distance (W_2) and Maximum Mean Discrepancy with a Gaussian kernel (MMD(G)).

Table 2: Trajectory inference results on the scRNA-seq dataset in mouse hematopoiesis and the single-cell qPCR iPSC dataset with bifurcation. All results are averages of 10 runs.

Methods	Mouse hematopoiesis				qPCR iPSC			
	W_2 (L)	MMD(L)	W_2 (F)	MMD(F)	W_2 (L)	MMD(L)	W_2 (F)	MMD(F)
TIGON	0.519	0.563	0.264	0.155	0.733	0.791	0.695	0.405
MIOFlow	0.514	0.629	0.220	0.056	0.770	1.039	0.345	0.155
OT-CFM	0.378	0.357	0.192	0.047	0.579	0.492	0.226	0.030
OT-DFM(C)	0.379	0.384	0.136	0.021	0.553	0.447	0.234	0.041
OT-DFM(D)	0.372	0.341	0.095	0.010	0.532	0.399	0.213	0.027

Mouse hematopoiesis dataset. We evaluate DFM on scRNA-seq data from mouse hematopoiesis (Weinreb et al., 2020), focusing on cells differentiating into neutrophils (Neu) and monocytes (Mo) at Day 2, 4, and 6. Models are trained on Day 2 and 6 data; after training, Day 2 samples are integrated forward to predict distributions at Days 4 and 6. Geodesic interpolation is used to construct trajectories. Due to the heterogeneity between pre- and post-differentiation distributions, trajectories inferred by ODE-based models (TIGON (Sha et al., 2024), MIOFlow (Huguet et al., 2022), OT-CFM) partially deviate from the data manifold, falling between Neu and Mo fates (Fig. 4 (b-e)). In contrast, OT-DFM(D), using distinct initial functions for Neu and Mo (Appendix D.3), preserves alignment with the data manifold (Fig. 4 (f)) and achieves more accurate predictions (Table 2).

qPCR iPSC dataset. We further apply DFM to model the bifurcation process of iPSCs in cardiomyocytes (Bargaje et al., 2017). Data from Day 2, 3, 4, and 5 are selected, with bifurcation observed from Day 3, where progenitor cells differentiate into mesodermal (M) and endodermal (En) fates. Models are trained on all time points except Day 3. After training, samples from Day 2 are forward-integrated to infer trajectories and predict distributions at Day 3 and 5. Due to distributional heterogeneity at the

bifurcation point, ODE-based methods and OT-DFM(C) produce trajectories that partially misalign with the data manifold (Fig. 4 (h–k)). In contrast, OT-DFM(D), by assigning separate initial functions for each fate (Appendix D.4), preserves trajectory alignment with the corresponding manifolds and yields more accurate predictions at both Day 3 and 5 compared to other methods (Table 2).

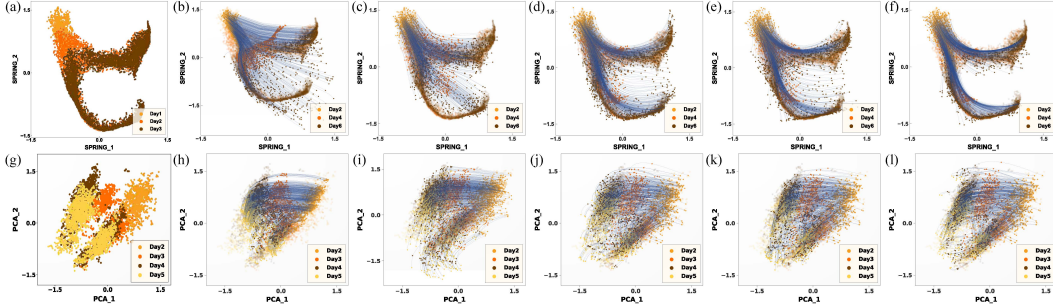


Figure 4: Comparisons between predicted trajectories of MIOFlow (b, h), TIGON (c, i), OT-CFM (d, j), OT-DFM(C) (e, k), and OT-DFM(D) (f, l) on scRNA-seq dataset in mouse hematopoiesis (a-f) and qPCR iPSC dataset (g-l). (a) and (g) illustrate the true data distribution.

6.4 IMAGE GENERATION

MNIST dataset. We design a *Semi-paired Image-to-Image Translation* task on the MNIST dataset, where the source domain consists of original images, and the target domain includes their negative images (the normalized pixel values transformed via $x \mapsto 1-x$). To provide partial supervision, 10% of the training data are paired with their negative counterparts as keypoints. During training, minibatches are independently sampled from the source and target distributions, and coupled via KPG-OT (KP-). The task aims to transform source images into their negatives, where, under linear interpolation, all transport paths intersect at an image with uniform pixel values of 0.5. While KP-CFM struggles with this transformation, KP-DFM(C) effectively handles it. As shown in Table 3, across varying delays τ , KP-DFM(C) consistently achieves lower FID scores than KP-CFM.

CIFAR-10 dataset. We further evaluate DFM on the CIFAR-10 dataset under a more general generation setup (Zhu & Lin, 2024), where the source distribution is a two-component Gaussian mixture. DFM employs trainable initial functions with distinct constant time derivatives, enabling generation from different mixture components to specific image classes. We compare CFM and DFM(D) under two coupling strategies: independent coupling (I-) and OT coupling (OT-). As shown in Table 4, I-CFM struggles with mode heterogeneity, while I-DFM(D) generates higher-quality images. With OT coupling, OT-DFM(D) consistently outperforms OT-CFM, especially when the number of function evaluations (NFE) is small. Additional results are provided in Appendix D.6.

Table 3: Comparison of FID between KP-CFM ($\tau = 0$) and KP-DFM(C) with different time delay on the MNIST dataset.

τ	0 (CFM)	0.125	0.250	0.500	1.000
FID	45.020	28.497	11.747	12.653	12.031

Table 4: Comparison of FID between CFM and DFM on the CIFAR-10 dataset.

NFE	10	20	30	40	Adap.
I-CFM	108.291	94.629	91.404	90.254	88.306
OT-CFM	78.165	27.512	16.409	12.026	6.162
I-DFM(D)	54.064	18.248	11.429	9.008	4.980
OT-DFM(D)	54.222	18.598	11.894	9.287	5.191

7 CONCLUSION

We introduce DFM, a novel continuous-time generative framework based on Neural DDEs. Through theoretical analysis, we highlight the limitations of ODE-based generative models, particularly their inability to capture certain transport strategies and preserve critical coupling information. In contrast, DFM offers universal approximation for arbitrary continuous transport strategies, addressing these shortcomings effectively. DFM also overcomes the challenge of transport between heterogeneous distributions by incorporating task-specific initial functions. Furthermore, it is naturally suited for modeling delay dynamical systems, a feature beyond the capability of ODEs. Extensive experiments on both synthetic and real-world datasets demonstrate that DFM achieves significantly more precise and versatile distribution transport strategies compared to FM.

ACKNOWLEDGMENTS

Q. Zhu is supported by the National Key R&D Program of China (No. 2025YFA1016503), by the National Natural Science Foundation of China (NSFC) (Nos. T2541024, 62406072, and 12171350), and by the STCSM (No. 23YF1402500). W. Lin is supported by the NSFC (Nos. 11925103, 12531018), the STCSM (Nos. 2021SHZDZX0103, 22JC1402500, 22JC1401402, 25JS2810400), and the SMEC (Nos. 2023ZKZD04, 2023KEJI05-72). X. Lu is supported by the National Science and Technology Major Project for Brain Science and Brain-like Intelligence Technology (No. 2025ZD0215700), the National Natural Science Foundation of China (Nos. 72025405, 72421002, 92467302, 72474223, 72301285), the Hunan Science and Technology Plan Project (Nos. 2023JJ40685, 2024RC3133), and the Major Program of Xiangjiang Laboratory (No. 24XJCYJ01001).

REFERENCES

- Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. *Arxiv Preprint Arxiv:2209.15571*, 2022.
- Uwe an der Heiden. Delays in physiological systems. *Journal of Mathematical Biology*, 8(4): 345–364, 1979.
- Rhishikesh Bargaje, Kalliopi Trachana, Martin N Shelton, Christopher S McGinnis, Joseph X Zhou, Cora Chadick, Savannah Cook, Christopher Cavanaugh, Sui Huang, and Leroy Hood. Cell population structure prior to bifurcation predicts efficiency of directed differentiation in human induced pluripotent cells. *Proceedings of the National Academy of Sciences*, 114(9):2271–2276, 2017.
- Charlotte Bunne, Geoffrey Schiebinger, Andreas Krause, Aviv Regev, and Marco Cuturi. Optimal transport for single-cell and spatial omics. *Nature Reviews Methods Primers*, 4(1):58, 2024.
- Sue Ann Campbell. Time delays in neural systems. In *Handbook of brain connectivity*, pp. 65–90. Springer, 2007.
- Yanne Chembo Kouomou, Pere Colet, Laurent Larger, and Nicolas Gastaud. Chaotic breathers in delayed electro-optical systems. *Physical Review Letters*, 95(20):203903, 2005.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in Neural Information Processing Systems*, 31, 2018.
- Yongxin Chen. *Modeling and control of collective dynamics: From Schrödinger bridges to optimal mass transport*. PhD thesis, University of Minnesota, 2016.
- Valentin De Bortoli, Guan-Horng Liu, Tianrong Chen, Evangelos A Theodorou, and Weilie Nie. Augmented bridge matching. *arXiv preprint arXiv:2311.06978*, 2023.
- Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. In *International Conference on Learning Representations*, 2022.
- Emilien Dupont, Arnaud Doucet, and Yee Whye Teh. Augmented neural odes. *Advances in Neural Information Processing Systems*, 32, 2019.
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *International Conference on Machine Learning*, 2024.
- TD Frank. Delay fokker-planck equations, novikov’s theorem, and boltzmann distributions as small delay approximations. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 72(1): 011112, 2005.
- David S Glass, Xiaofan Jin, and Ingmar H Riedel-Kruse. Nonlinear delay differential equations and their application to modeling biological network motifs. *Nature Communications*, 12(1):1788, 2021.

- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27, 2014.
- Will Grathwohl, Ricky TQ Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. Ffjord: Free-form continuous dynamics for scalable reversible generative models. *Arxiv Preprint Arxiv:1810.01367*, 2018.
- Xiang Gu, Yucheng Yang, Wei Zeng, Jian Sun, and Zongben Xu. Keypoint-guided optimal transport with applications in heterogeneous domain adaptation. *Advances in Neural Information Processing Systems*, 35:14972–14985, 2022.
- Xiang Gu, Liwei Yang, Jian Sun, and Zongben Xu. Optimal transport-guided conditional score-based diffusion model. *Advances in Neural Information Processing Systems*, 36:36540–36552, 2023.
- Steve Guillouezic, Ivan L’Heureux, and André Longtin. Small delay approximation of stochastic delay differential equations. *Physical Review E*, 59(4):3970, 1999.
- Pengsheng Guo and Alex Schwing. Variational rectified flow matching. In *Forty-second International Conference on Machine Learning*, 2025.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- Guillaume Hugué, Daniel Sumner Magruder, Alexander Tong, Oluwadamilola Fasina, Manik Kuchroo, Guy Wolf, and Smita Krishnaswamy. Manifold interpolating optimal-transport flows for trajectory inference. *Advances in neural information processing systems*, 35:29705–29718, 2022.
- Xunbi A. Ji and Gábor Orosz. Trainable delays in time delay neural networks for learning delayed dynamics. *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–11, 2024. doi: 10.1109/TNNLS.2024.3379020.
- Leonid V Kantorovich. On the translocation of masses. In *Dokl. Akad. Nauk. USSR (NS)*, volume 37, pp. 199–201, 1942.
- Kacper Kapuśniak, Peter Potapchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael Bronstein, Avishek Joey Bose, and Francesco Di Giovanni. Metric flow matching for smooth interpolations on the data manifold. *arXiv preprint arXiv:2405.14780*, 2024.
- Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- Xingchao Liu, Chengyue Gong, et al. Flow straight and fast: learning to generate and transfer data with rectified flow. In *International Conference on Learning Representations*, 2023.
- AJ Lotka. Elements of physical biology. *Williams and Wilkins*, 1925.
- Huanfei Ma, Bing Xu, Wei Lin, and Jianfeng Feng. Adaptive identification of time delays in nonlinear dynamical models. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 82(6): 066210, 2010.
- Huanfei Ma, Siyang Leng, Chenyang Tao, Xiong Ying, Jürgen Kurths, Ying-Cheng Lai, and Wei Lin. Detection of time delays and directional interactions based on time series from complex dynamical systems. *Physical Review E*, 96(1):012221, 2017.
- Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021.
- George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021.

- Dogyun Park, Sojin Lee, Sihyeon Kim, Taehoon Lee, Youngjoon Hong, and Hyunwoo J Kim. Constant acceleration flow. *Advances in Neural Information Processing Systems*, 37:90030–90060, 2024.
- Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman, and Ricky TQ Chen. Multisample flow matching: Straightening flows with minibatch couplings. *arXiv preprint arXiv:2304.14772*, 2023.
- Aram-Alexandre Pooladian, Carles Domingo-Enrich, Ricky TQ Chen, and Brandon Amos. Neural optimal transport with lagrangian costs. In *Uncertainty in Artificial Intelligence*, pp. 2989–3003. PMLR, 2024.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*, pp. 1278–1286. PMLR, 2014.
- Fathalla A Rihan et al. Delay differential equations and applications to biology. 2021.
- Wouter Saelens, Robrecht Cannoodt, Helena Todorov, and Yvan Saeys. A comparison of single-cell trajectory inference methods. *Nature Biotechnology*, 37(5):547–554, 2019.
- Benjamin Sanchez-Lengeling and Alán Aspuru-Guzik. Inverse molecular design using machine learning: Generative models for matter engineering. *Science*, 361(6400):360–365, 2018.
- Filippo Santambrogio. Optimal transport for applied mathematicians. *Birkäuser, NY*, 55(58-63):94, 2015.
- Yutong Sha, Yuchi Qiu, Peijie Zhou, and Qing Nie. Reconstructing growth and dynamic trajectories from single-cell transcriptomics data. *Nature Machine Intelligence*, 6(1):25–39, 2024.
- Sagar Shrestha and Xiao Fu. Diversified flow matching with translation identifiability. In *Forty-second International Conference on Machine Learning*, 2025.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pp. 2256–2265. PMLR, 2015.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *Arxiv Preprint Arxiv:2010.02502*, 2020a.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020b.
- Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in neural information processing systems*, 34:1415–1428, 2021.
- Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguët, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *arXiv preprint arXiv:2302.00482*, 2023a.
- Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguët, Guy Wolf, and Yoshua Bengio. Simulation-free schrödinger bridges via score and flow matching. *arXiv preprint arXiv:2307.03672*, 2023b.
- Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.

- Caleb Weinreb, Alejo Rodriguez-Fraticelli, Fernando D Camargo, and Allon M Klein. Lineage tracing on transcriptional landscapes links state to fate during differentiation. *Science*, 367(6479): eaaw3381, 2020.
- 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. *ACM Computing Surveys*, 56(4):1–39, 2023.
- Yichi Zhang, Yici Yan, Alex Schwing, and Zhizhen Zhao. Towards hierarchical rectified flow. In *The Thirteenth International Conference on Learning Representations*, 2025.
- Qunxi Zhu and Wei Lin. Switched flow matching: Eliminating singularities via switching odes. In *Forty-first International Conference on Machine Learning*, 2024.
- Qunxi Zhu, Yao Guo, and Wei Lin. Neural delay differential equations. In *International Conference on Learning Representations*, 2020.
- Qunxi Zhu, Yifei Shen, Dongsheng Li, and Wei Lin. Neural piecewise-constant delay differential equations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 9242–9250, 2022.

The structure of the appendix is as follows:

- Appendix A provides the connection between the probability flows of ODEs and DDEs.
- Appendix B provides the formal proofs for the theoretical results presented in the main text.
- Appendix C provides pseudocode for DFM(C) and DFM(D).
- Appendix D includes the experimental setup details and additional experimental results.
- Appendix E provides the comparison of the training and inference time required by DFM and CFM.
- Appendix F provides the sensitivity analysis of time delay parameter τ on both synthetic and real-world datasets.
- Appendix G compares DFM with other ODE-based generative methods designed to address trajectory intersections and distributional heterogeneity, further highlighting the advantages of the DDE-based framework.

A CONNECTION BETWEEN THE PROBABILITY FLOWS OF ODES AND DDES

Based on the vector field $\mathbf{u}[t, \mathbf{x}(t), \mathbf{x}(t - \tau)]$ in Eq. (1), we can define the conditional average drift (CAD) without delay terms as follows (Guillouzic et al., 1999):

$$\bar{\mathbf{u}}(t, \mathbf{x}|\boldsymbol{\psi}) = \int d\mathbf{x}_\tau [\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau) p(\mathbf{x}_\tau, t - \tau | \mathbf{x}, t; \boldsymbol{\psi})], \quad (24)$$

where $p(\mathbf{x}_\tau, t - \tau | \mathbf{x}, t; \boldsymbol{\psi})$ denotes the conditional probability given $\mathbf{x}(t) = \mathbf{x}$. Note that $p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \boldsymbol{\psi}) = p(\mathbf{x}_\tau, t - \tau | \mathbf{x}, t; \boldsymbol{\psi}) p(\mathbf{x}, t | \boldsymbol{\psi})$, so Eq. (2) is equivalent to:

$$\frac{\partial p(\mathbf{x}, t | \boldsymbol{\psi})}{\partial t} = -\nabla \cdot \{ \bar{\mathbf{u}}(t, \mathbf{x} | \boldsymbol{\psi}) p(\mathbf{x}, t | \boldsymbol{\psi}) \}, \quad (25)$$

which is precisely the continuity equation satisfied by the probability flows of ODEs.

Remark A.1. This implies that a non-delayed vector field can be constructed to match the probability flow of the delayed vector field in Eq. (1). However, the integration in Eq. (24) eliminates the coupling information between the states at times t and $t - \tau$. Consequently, while the ODE preserves the marginal probability flow, it fails to maintain the coupling relationship during distribution transfer, as illustrated by a simple example in Fig. 1 (a).

B PROOFS OF THEORETICAL RESULTS

B.1 FOKKER-PLANCK EQUATION OF DDES

Theorem B.1 (Fokker-Planck equation of DDEs, (Guillouzic et al., 1999; Frank, 2005)). *Consider a time-dependent DDE with a single delay term:*

$$\begin{aligned} \frac{d\mathbf{x}(t)}{dt} &= \mathbf{u}[t, \mathbf{x}(t), \mathbf{x}(t - \tau)], \quad t \in [0, T], \\ \mathbf{x}(h) &= \boldsymbol{\psi}(h), \quad h \in [-\tau, 0], \end{aligned} \quad (26)$$

where $\mathbf{u}[t, \mathbf{x}(t), \mathbf{x}(t - \tau)] : [0, T] \times \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a smooth vector field which is abbreviated as $\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau)$ in the following, $\boldsymbol{\psi}(h)$ represents the continuous initial function. The associated probability flows $p(\mathbf{x}, t | \boldsymbol{\psi}) : \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}^+$ satisfy the delay Fokker-Planck equation:

$$\frac{\partial p(\mathbf{x}, t | \boldsymbol{\psi})}{\partial t} = -\nabla \cdot \left\{ \int d\mathbf{x}_\tau [\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau) p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \boldsymbol{\psi})] \right\}, \quad (27)$$

where $p(\mathbf{x}, t | \boldsymbol{\psi})$ is the probability density at time t given the initial function, while $p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \boldsymbol{\psi})$ is the joint probability density representing the likelihood of the system being at \mathbf{x} at time t and at \mathbf{x}_τ at time $t - \tau$.

Proof. The proof follows a similar approach as presented in Guillouzic et al. (1999). Without loss of generality, we assume that $d = 1$ and the state variable $x \in [a, b]$ in Eq. (26). Consider an arbitrary C^2 function $F(x)$ defined on the interval $[a, b]$, i.e. $F \in C^2([a, b])$, which satisfies the following conditions:

$$\lim_{x \rightarrow a} F(x) = \lim_{x \rightarrow b} F(x) = 0, \quad (28)$$

$$\lim_{x \rightarrow a} \frac{d}{dx} F(x) = \lim_{x \rightarrow b} \frac{d}{dx} F(x) = 0. \quad (29)$$

Then, by applying the Taylor expansion, we obtain:

$$dF[x(t)] = F[x(t) + dx(t)] - F[x(t)] = \left\{ u[t, x(t), x(t - \tau)] \frac{d}{dx} F[x(t)] \right\} dt. \quad (30)$$

The ensemble average (average over realizations) of $dF[x(t)]$ can be written as

$$\left\langle \frac{d}{dt} F[x(t)] \right\rangle = \left\langle u[t, x(t), x(t - \tau)] \frac{d}{dx} F[x(t)] \right\rangle. \quad (31)$$

We denote $p(x, t; x_\tau, t - \tau | \psi) dx dx_\tau$ as the probability that $x(t) \in [x, x + dx]$ and $x(t - \tau) \in [x_\tau, x_\tau + dx_\tau]$ given the initial function ψ . Then, Eq. (31) is equivalent to

$$\begin{aligned} & \int_a^b dx F(x) \int_a^b dx_\tau \frac{\partial}{\partial t} p(x, t; x_\tau, t - \tau | \psi) \\ &= \int_a^b dx \frac{d}{dx} F[x(t)] \int_a^b dx_\tau u(t, x, x_\tau) p(x, t; x_\tau, t - \tau | \psi). \end{aligned} \quad (32)$$

By applying the integration by parts formula to the right-hand side, we obtain:

$$\begin{aligned} & \int_a^b dx F(x) \int_a^b dx_\tau \frac{\partial}{\partial t} p(x, t; x_\tau, t - \tau | \psi) \\ &= \int_a^b dx F(x) \int_a^b dx_\tau \left\{ -\frac{\partial}{\partial x} [u(t, x, x_\tau) p(x, t; x_\tau, t - \tau | \psi)] \right\}, \end{aligned} \quad (33)$$

where the surface terms are neglected due to Eq. (28) and Eq. (29). Since $F(x)$ is arbitrary, Eq. (33) leads to

$$\frac{\partial}{\partial t} p(x, t | \psi) = -\frac{\partial}{\partial x} \left\{ \int_a^b dx_\tau [u(t, x, x_\tau) p(x, t; x_\tau, t - \tau | \psi)] \right\}. \quad (34)$$

□

B.2 PROOFS OF PROPOSITIONS IN THE MAIN TEXT

Proposition 3.1 (Restriction on trajectory intersections). *Suppose that the target transport strategy, guided by certain key points or specific coupling constraints, inevitably induces trajectory intersections in the augmented phase space (t, \mathbf{x}) during the transport process. If the ODE-based FM is governed by a vector field that is Lipschitz continuous in \mathbf{x} and continuous in t , then the corresponding flow map cannot precisely preserve the transport strategy.*

Proof of Proposition 3.1. Consider a time-dependent Neural ODE corresponding to a ODE-based FM with the following form:

$$\begin{aligned} \frac{d\mathbf{x}(t)}{dt} &= \mathbf{v}[t, \mathbf{x}(t); \theta], \quad t \in [0, T], \\ \mathbf{x}(0) &= \mathbf{x}_0, \end{aligned} \quad (35)$$

where $\mathbf{v}[t, \mathbf{x}(t); \theta] : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a parameterized vector field. Since the vector field is Lipschitz continuous in \mathbf{x} and continuous in t , by the Picard-Lindelöf Theorem, the corresponding initial value problem (35) has a unique solution on the interval $[0, T]$. Suppose that there exist two distinct solutions $\mathbf{x}^1(t)$ and $\mathbf{x}^2(t)$ corresponding to different initial values \mathbf{x}_0^1 and \mathbf{x}_0^2 , which intersect at (t^*, \mathbf{x}^*) . By the uniqueness of the solution, these two trajectories must correspond to the same solution, leading to a contradiction. Therefore, given a transport strategy, which inevitably leads to trajectory intersections during the transport process, the flow map corresponding to the ODE-based FM cannot precisely ensure the transport strategy. □

Proposition 3.3 (Heterogeneity in distributions). *Assume that the source (resp., target) distribution q_0 (resp., q_1) is supported on M (resp., N) disjoint compact sets $\{U_i^0\}_{i=1}^M$ (resp., $\{U_j^1\}_{j=1}^N$), where each $U_i^0 \in \mathbb{R}^d$ (resp., $U_j^1 \in \mathbb{R}^d$) is a path-connected set with non-zero measure and $q_0(\mathbf{x}_0) > 0$ (resp., $q_1(\mathbf{x}_1) > 0$) for all $\mathbf{x}_0 \in \bigcup_{i=1}^M U_i^0$ (resp., $\mathbf{x}_1 \in \bigcup_{j=1}^N U_j^1$). If $M < N$, then for any transport map $T : \bigcup_{i=1}^M U_i^0 \rightarrow \bigcup_{j=1}^N U_j^1$ satisfying $T_{\#}q_0 = q_1$, the following holds: if the flow map induced by a Neural ODE is equal to the transport map T , then the associated vector field cannot be Lipschitz continuous.*

Proof of Proposition 3.3. Assume there exists a parameterized vector field $\mathbf{v}(t, \mathbf{x}; \theta)$ whose flow map ϕ_T coincides with the transport map T , i.e., $\phi_T(\mathbf{x}) = T(\mathbf{x})$ for all $\mathbf{x} \in \bigcup_{i=1}^M U_i^0$.

First, consider the case $M < N$. Then there must exist $m \in \{1, 2, \dots, M\}$, and $n_1, n_2 \in \{1, 2, \dots, N\}$ such that $q_1[T(U_m^0) \cap U_{n_1}^1] > 0$ and $q_1[T(U_m^0) \cap U_{n_2}^1] > 0$. Otherwise, for each $m \in \{1, 2, \dots, M\}$, there exists a unique $n \in \{1, 2, \dots, N\}$ such that $q_1[T(U_m^0) \cap U_n^1] = 0$ for all $n' \neq n$. Since $M < N$, there exists $n^* \in \{1, 2, \dots, N\}$ such that $\sum_{i=1}^M q_1[T(U_i^0) \cap U_{n^*}^1] = 0$, which implies

$$q_1(U_{n^*}^1) = q_1 \left[T \left(\bigcup_{i=1}^M U_i^0 \right) \cap U_{n^*}^1 \right] = q_1 \left[\bigcup_{i=1}^M T(U_i^0) \cap U_{n^*}^1 \right] = 0, \quad (36)$$

a contradiction.

Hence, we can choose $\mathbf{x}_{m,1}, \mathbf{x}_{m,2} \in U_m^0$ such that $T(\mathbf{x}_{m,1}) \in U_{n_1}^1$, $T(\mathbf{x}_{m,2}) \in U_{n_2}^1$. By the path-connectedness of U_m^0 , there exists a continuous path $\gamma : [0, 1] \rightarrow U_m^0$ with $\gamma(0) = \mathbf{x}_{m,1}$ and $\gamma(1) = \mathbf{x}_{m,2}$. Then, for any $\epsilon > 0$, there exist $\mathbf{x}_s, \tilde{\mathbf{x}}_s \in \gamma$ such that $|\mathbf{x}_s - \tilde{\mathbf{x}}_s| < \epsilon$, and $T(\mathbf{x}_s)$ and $T(\tilde{\mathbf{x}}_s)$ lie in distinct subsets of $\{U_j^1\}_{j=1}^N$. Since each U_j^1 is a compact subset of \mathbb{R}^d , it follows that for any $j_1, j_2 \in \{1, 2, \dots, N\}$, $j_1 \neq j_2$,

$$\text{dist}(U_{j_1}^1, U_{j_2}^1) = \inf \{ |\mathbf{x} - \mathbf{y}| \mid \mathbf{x} \in U_{j_1}^1, \mathbf{y} \in U_{j_2}^1 \} > 0. \quad (37)$$

Let

$$d = \min_{j_1 \neq j_2} \{ \text{dist}(U_{j_1}^1, U_{j_2}^1) \}_{j_1, j_2} > 0. \quad (38)$$

Then we have

$$\frac{|\phi_T(\mathbf{x}_s) - \phi_T(\tilde{\mathbf{x}}_s)|}{|\mathbf{x}_s - \tilde{\mathbf{x}}_s|} = \frac{|T(\mathbf{x}_s) - T(\tilde{\mathbf{x}}_s)|}{|\mathbf{x}_s - \tilde{\mathbf{x}}_s|} > \frac{d}{\epsilon}. \quad (39)$$

On the other hand, if $\mathbf{v}(t, \mathbf{x}; \theta)$ is Lipschitz continuous with constant $L \geq 0$, then for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$,

$$|\mathbf{v}(t, \mathbf{x}; \theta) - \mathbf{v}(t, \mathbf{y}; \theta)| \leq L|\mathbf{x} - \mathbf{y}|, \quad (40)$$

By Grönwall's inequality, the corresponding flow satisfies

$$\begin{aligned} |\phi_t(\mathbf{x}) - \phi_t(\mathbf{y})| &= \left| \mathbf{x} - \mathbf{y} + \int_0^t \mathbf{v}[s, \mathbf{x}(s); \theta] - \mathbf{v}[s, \mathbf{y}(s); \theta] ds \right| \\ &\leq |\mathbf{x} - \mathbf{y}| + \int_0^t |\mathbf{v}[s, \mathbf{x}(s); \theta] - \mathbf{v}[s, \mathbf{y}(s); \theta]| ds \\ &\leq |\mathbf{x} - \mathbf{y}| + \int_0^t L|\phi_s(\mathbf{x}) - \phi_s(\mathbf{y})| ds. \end{aligned} \quad (41)$$

Thus,

$$|\phi_T(\mathbf{x}) - \phi_T(\mathbf{y})| \leq e^{LT} |\mathbf{x} - \mathbf{y}|, \quad (42)$$

holds for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$. Now, choose $\epsilon < \frac{d}{e^{LT}}$. Then from (39), there exist $\mathbf{x}_s, \tilde{\mathbf{x}}_s \in U_m^0$ such that

$$|\phi_T(\mathbf{x}_s) - \phi_T(\tilde{\mathbf{x}}_s)| > \frac{d}{\epsilon} |\mathbf{x}_s - \tilde{\mathbf{x}}_s| > e^{LT} |\mathbf{x}_s - \tilde{\mathbf{x}}_s|, \quad (43)$$

which contradicts the Lipschitz bound.

Therefore, any parameterized vector field realizing such a transport map cannot be Lipschitz continuous. This completes the proof. \square

Proposition 4.1. *The marginal vector field $\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau)$ given in Eq. (15), together with the selected initial function ψ , generates the probability path $p(\mathbf{x}, t|\psi)$ and the joint probability path $p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi)$ in Eq. (14).*

Proof of Proposition 4.1. Given that $\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau|z)$ is the conditional vector field that generates $p(\mathbf{x}, t|z, \psi)$ and $p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|z, \psi)$, it is evident that the following delay Fokker-Planck equation holds:

$$\frac{\partial p(\mathbf{x}, t|z, \psi)}{\partial t} = -\nabla \cdot \left\{ \int d\mathbf{x}_\tau [\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau|z) p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|z, \psi)] \right\}. \quad (44)$$

To verify this proposition, we only need to check that the marginal vector field $\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau)$ given by Eq. (15), the marginal probability density path $p(\mathbf{x}, t|\psi)$ and the joint probability density $p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi)$ satisfy the following delay Fokker-Planck equation:

$$\frac{\partial p(\mathbf{x}, t|\psi)}{\partial t} = -\nabla \cdot \left\{ \int d\mathbf{x}_\tau [\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau) p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi)] \right\}. \quad (45)$$

Assuming that all the functions involved satisfy the regularity conditions necessary for the interchange of integration and differentiation, we have that

$$\begin{aligned} & \frac{\partial p(\mathbf{x}, t|\psi)}{\partial t} \\ &= \frac{\partial}{\partial t} \int p(\mathbf{x}, t|z, \psi) q(z) dz \\ &= \int \left[\frac{\partial}{\partial t} p(\mathbf{x}, t|z, \psi) \right] q(z) dz \\ &= - \int \left(\nabla \cdot \left\{ \int d\mathbf{x}_\tau [\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau|z) p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|z, \psi)] \right\} \right) q(z) dz \\ &= -\nabla \cdot \left\{ \int d\mathbf{x}_\tau \left[\int \mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau|z) p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|z, \psi) q(z) dz \right] \right\} \\ &= -\nabla \cdot \left\{ \int d\mathbf{x}_\tau \left[\int \frac{\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau|z) p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|z, \psi)}{p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi)} q(z) dz \cdot p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi) \right] \right\} \\ &= -\nabla \cdot \left\{ \int d\mathbf{x}_\tau [\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau) p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi)] \right\}. \end{aligned} \quad (46)$$

□

Proposition 4.2. *Given that $p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi) > 0$ for all $\mathbf{x}, \mathbf{x}_\tau \in \mathbb{R}^d$ and $t \in [0, 1]$, up to a constant independent of θ , $\mathcal{L}_{DCFM}(\theta)$ and $\mathcal{L}_{DFM}(\theta)$ are equal, which further implies that $\nabla_\theta \mathcal{L}_{DCFM}(\theta) = \nabla_\theta \mathcal{L}_{DFM}(\theta)$.*

Proof of Proposition 4.2. To guarantee the existence of all integrals and the validity of changing the order of integration (as justified by Fubini's Theorem), we assume that $p(\mathbf{x}, t|z, \psi)$, $p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|z, \psi)$ decay to zero sufficiently fast as $\|\mathbf{x}\| \rightarrow \infty$ and $\|\mathbf{x}_\tau\| \rightarrow \infty$, and further assume that $\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau|z)$, $\mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta)$, $\nabla_\theta \mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta)$ are bounded. Note that $t \sim \mathcal{U}(0, 1)$ and $\psi \sim q^\circ(\psi)$ are both independent of θ , so we fixed t and ψ in the following analysis. Using the bilinearity of the Euclidean norm and the independence of $\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau|z)$ from θ , we have:

$$\begin{aligned} & \nabla_\theta \mathbb{E}_{p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi)} \|\mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta) - \mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau)\|^2 \\ &= \nabla_\theta \mathbb{E}_{p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi)} (\|\mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta)\|^2 - 2 \langle \mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta), \mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau) \rangle + \|\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau)\|^2) \\ &= \nabla_\theta \mathbb{E}_{p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|\psi)} (\|\mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta)\|^2 - 2 \langle \mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta), \mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau) \rangle), \end{aligned} \quad (47)$$

and

$$\begin{aligned} & \nabla_\theta \mathbb{E}_{q(z), p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|z, \psi)} \|\mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta) - \mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau|z)\|^2 \\ &= \nabla_\theta \mathbb{E}_{q(z), p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|z, \psi)} (\|\mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta)\|^2 - 2 \langle \mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta), \mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau|z) \rangle + \|\mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau|z)\|^2) \\ &= \nabla_\theta \mathbb{E}_{q(z), p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau|z, \psi)} (\|\mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta)\|^2 - 2 \langle \mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta), \mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau|z) \rangle). \end{aligned} \quad (48)$$

Next,

$$\begin{aligned}
& \mathbb{E}_{p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \boldsymbol{\psi})} \|\mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta)\|^2 \\
&= \iint \|\mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta)\|^2 p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \boldsymbol{\psi}) d\mathbf{x} d\mathbf{x}_\tau \\
&= \iiint \|\mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta)\|^2 p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \mathbf{z}, \boldsymbol{\psi}) q(\mathbf{z}) d\mathbf{z} d\mathbf{x} d\mathbf{x}_\tau \\
&= \mathbb{E}_{q(\mathbf{z}), p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \mathbf{z}, \boldsymbol{\psi})} \|\mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta)\|^2.
\end{aligned} \tag{49}$$

Finally,

$$\begin{aligned}
& \mathbb{E}_{p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \boldsymbol{\psi})} \langle \mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta), \mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau) \rangle \\
&= \iint \left\langle \mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta), \frac{\int \mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau | \mathbf{z}) p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \mathbf{z}, \boldsymbol{\psi}) q(\mathbf{z}) d\mathbf{z}}{p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \boldsymbol{\psi})} \right\rangle p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \boldsymbol{\psi}) d\mathbf{x} d\mathbf{x}_\tau \\
&= \iint \left\langle \mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta), \int \mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau | \mathbf{z}) p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \mathbf{z}, \boldsymbol{\psi}) q(\mathbf{z}) d\mathbf{z} \right\rangle d\mathbf{x} d\mathbf{x}_\tau \\
&= \iiint \langle \mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta), \mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau | \mathbf{z}) \rangle p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \mathbf{z}, \boldsymbol{\psi}) q(\mathbf{z}) d\mathbf{z} d\mathbf{x} d\mathbf{x}_\tau \\
&= \mathbb{E}_{q(\mathbf{z}), p(\mathbf{x}, t; \mathbf{x}_\tau, t - \tau | \mathbf{z}, \boldsymbol{\psi})} \langle \mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta), \mathbf{u}(t, \mathbf{x}, \mathbf{x}_\tau | \mathbf{z}) \rangle.
\end{aligned} \tag{50}$$

From Eq. (49) and Eq. (50), it follows that Eq. (47) is equal to Eq. (48) for any latent variable \mathbf{z} , which can be further deduced that $\nabla_{\theta} \mathcal{L}_{\text{DCFM}}(\theta) = \nabla_{\theta} \mathcal{L}_{\text{DFM}}(\theta)$. \square

Proposition 4.3 (Universal approximating capability of DDE’s flow map). *For any given $\epsilon > 0$ and continuous transport map $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$, which push-forward the source distribution q_0 to the target distribution q_1 , i.e. $F_{\#} q_0 = q_1$, if there exists a neural network $\mathbf{f}(\mathbf{x}; \theta)$ such that $\|\mathbf{f}(\mathbf{x}; \theta) - [F(\mathbf{x}) - \mathbf{x}]\| < \epsilon$, for all $\mathbf{x} \in \mathbb{R}^d$, then we can construct a vector field with a single delay term $\mathbf{v}(t, \mathbf{x}, \mathbf{x}_\tau; \theta)$, where the corresponding flow map $G(\mathbf{x}; \theta)$, under the constant initial function condition, satisfies $\|G(\mathbf{x}; \theta) - F(\mathbf{x})\| < \epsilon$, for all $\mathbf{x} \in \mathbb{R}^d$.*

Proof of Proposition 4.3. The proof is straightforward. We consider the following DDE with a constant initial function and a parameterized vector field with a single delay term:

$$\begin{aligned}
\frac{d\mathbf{x}(t)}{dt} &= \mathbf{v}[\mathbf{x}(t - \tau); \theta], \quad t \in [0, 1], \\
\mathbf{x}(t) &\equiv \mathbf{x}_0, \quad t \in [-\tau, 0],
\end{aligned} \tag{51}$$

where the time delay $\tau = 1$, and the vector field depends solely on the delay term, independent of the current state and time, which is a special degenerate case of DFM. In this case, for any point sampled from the initial distribution $\mathbf{x}_0 \sim q_0$, the corresponding vector field $\mathbf{v}[\mathbf{x}(t - \tau); \theta] = \mathbf{v}(\mathbf{x}_0; \theta)$ remains constant over the time interval $t \in [0, 1]$, which means that the flow map $G : \mathbb{R}^d \rightarrow \mathbb{R}^d$ associated with the above DDE will map \mathbf{x}_0 to $\mathbf{x}_0 + \mathbf{v}[\mathbf{x}_0; \theta] \cdot 1$. For any given $\epsilon > 0$, suppose there exists a neural network $\mathbf{f}(\mathbf{x}; \theta)$ such that $\|\mathbf{f}(\mathbf{x}; \theta) - [F(\mathbf{x}) - \mathbf{x}]\| < \epsilon$, for all $\mathbf{x} \in \mathbb{R}^d$. Let $\mathbf{v}(\mathbf{x}_0; \theta) = \mathbf{f}(\mathbf{x}_0; \theta)$. Then it follows that $\|G(\mathbf{x}_0; \theta) - F(\mathbf{x}_0)\| = \|\mathbf{x}_0 + \mathbf{v}(\mathbf{x}_0; \theta) \cdot 1 - F(\mathbf{x}_0)\| \leq \|\mathbf{v}(\mathbf{x}_0; \theta) - \mathbf{f}(\mathbf{x}_0; \theta)\| + \|\mathbf{f}(\mathbf{x}_0; \theta) - [F(\mathbf{x}_0) - \mathbf{x}_0]\| < \epsilon$. This implies that the flow map of the Neural DDEs approximately preserves the target transport strategy, and naturally, it can approximately push-forward q_0 to q_1 . \square

C ALGORITHMS

To more clearly illustrate the overall workflow of our proposed DFM, we provide pseudocode for different DFM methods: Algorithm 1 for DFM(C) and Algorithm 2 for DFM(D).

D EXPERIMENTAL SETUP DETAILS AND ADDITIONAL RESULTS

In this section, we provide a detailed explanation of the experimental settings for different datasets described in the main text and present additional experimental results. The experimental details for the delay dynamical systems and single-cell datasets are summarized in Table 5).

Algorithm 1 Delay Flow Matching with Constant Initial Functions

```

# Input: Data={X0,X1}, the source data X0 ~ q0 and the target data
X1 ~ q1; Time delay τ
# Output: Model v(t,x,xτ;θ) for the vector field with a single
delay term xτ = x(t-τ)
Initialize Model: v(t,x,xτ;θ) # Arbitrary neural network structure
Select Coupling Strategy: KPG-OT/OT/Independent
Select Interpolation Method: Linear/Geodesic/CSpline # Consider
the linear interpolation below
Initialize Initial Function: ψ*(t;x0) ≡ x0 for t ∈ [-τ,0]
While not converge:
    Sample x0 ⊂ X0, x1 ⊂ X1 # Generate batch data
    Optimizer.zero_grad()
    x0,x1 = Coupling-Strategy(x0, x1) # Construct coupling
    t = torch.rand(batchsize) # Randomly sample t ∈ [0,1]
    x(t) = (1-t) × x0 + t × x1 # Sample state at t based on the linear
interpolation
    x(t-τ) = (1-t+τ) × x0 + (t-τ) × x1 for t ∈ [τ,1] # Sample delay state
at t-τ based on the linear interpolation
    x(t-τ) = x0 for t ∈ [0,τ) # Sample delay state at t-τ based on
the constant initial function
    Loss = { v[t,x(t),x(t-τ);θ] - (x1-x0) }.pow(2).mean()
    Loss.backward()
    Optimizer.step()
return Model

```

Table 5: The experimental setup details for delay dynamical systems and single-cell datasets.

Setup		Autoregulation	Spiral DDE	Mouse hematopoiesis	iPSCs
Data	Dimension	1	2	2	4
Structure	Hidden layer	3	3	3	3
	Hidden neuron	64	64	64	64
	Activation	Tanh	Tanh	SELU	SELU
	Time input	False	False	True	True
	Delay term input	True (DFM) / False (CFM)	True (DFM) / False (CFM)	True (DFM) / False (CFM)	True (DFM) / False (CFM)
Training	Coupling	KPG-OT	OT	OT	OT
	Initial function (DFM)	Constant	Constant	Diverse	Diverse
	Latent variable γ	CSpline	CSpline	Geodesic	CSpline
	Batch size	256	256	128	128
	Iteration	2k	2k	10k	2k
	Optimizer	Adam	Adam	Adam	Adam
Learning rate	10 ⁻³	10 ⁻³	10 ⁻³	10 ⁻³	
Results	Figures & tables	Figs. 3 & Tab. 6, 7, 8	Figs. 3 & Tab.9, 10	Fig. 4 & Tab. 2	Fig. 4 & Tab. 2

D.1 BIOLOGICAL AUTOREGULATION MOTIF

Dataset generation. The autoregulation motif, which is one of the most common biological network motifs, can be modeled as the following 1-d DDE:

$$\dot{x}(t) = \frac{\eta}{1 + x^n(t - \tau)}, \quad (52)$$

where $\tau = 1$, $n = 2$, $\eta = 5$. As shown in Fig. 3, the dynamics exhibit damped oscillations. We sample 1,000 initial points from the uniform distribution $\mathcal{U}(0.2, 1.2)$ and perform forward integration based on constant initial functions $\psi(h; \mathbf{x}_0) = \mathbf{x}_0$, $h \in [-1, 0]$, to obtain trajectories. After that, we select snapshots at $t = 0, 1, 2, \dots, 7$ as the training dataset.

Details for training. For each training iteration, we randomly and independently sample 256 data points from each snapshot. Pairing between minibatch data points from adjacent time steps is performed using KPG-OT. For this task, we utilize the true coupling of only two keypoint pairs between adjacent time steps as guidance, although incorporating more keypoint pairs would enable more accurate pairing. Once data points across time steps are fully matched, we construct 256

Algorithm 2 Delay Flow Matching with Diverse Initial Functions

```

# Input: Data={ $\mathbf{X}_0, \mathbf{X}_1$ }, the source data  $\mathbf{X}_0 \sim q_0$  and the target data  $\mathbf{X}_1 \sim q_1$ ; Time delay  $\tau$ 
# Output: Model  $v(t, \mathbf{x}, \mathbf{x}_\tau; \theta)$  for the vector field with a single delay term  $\mathbf{x}_\tau = \mathbf{x}(t - \tau)$ 
Initialize Model:  $v(t, \mathbf{x}, \mathbf{x}_\tau; \theta)$  # Arbitrary neural network structure
Select Coupling Strategy: KPG-OT/OT/Independent
Select Interpolation Method: Linear/Geodesic/CSpline # Consider the linear interpolation below
Select Clustering Method: Gaussian Mixture Model/DBSCAN
Cluster the source and target dataset:
 $\mathbf{X}_0^{(1)}, \dots, \mathbf{X}_0^{(M)} = \text{Clustering\_Method}(\mathbf{X}_0)$ 
 $\mathbf{X}_1^{(1)}, \dots, \mathbf{X}_1^{(N)} = \text{Clustering\_Method}(\mathbf{X}_1)$ 
Initialize Diverse Initial Functions for Different Clusters:
  Select different constants  $C_{mn}$ , for  $\mathbf{x}_0 \in X_0^{(m)}$ ,  $\mathbf{x}_1 \in X_1^{(n)}$ :
     $\psi_{mn}^*(t; \mathbf{x}_0) = \mathbf{x}_0 + C_{mn} \times t$  for  $t \in [-\tau, 0]$ 
While not converge:
  Sample  $\mathbf{x}_0 \subset X_0$ ,  $\mathbf{x}_1 \subset X_1$  # Generate batch data
  Optimizer.zero_grad()
   $\mathbf{x}_0, \mathbf{x}_1 = \text{Coupling\_Strategy}(\mathbf{x}_0, \mathbf{x}_1)$  # Construct coupling
  Identify the cluster membership of  $\mathbf{x}_0$ :  $I_{\mathbf{x}_0}$ 
  Identify the cluster membership of  $\mathbf{x}_1$ :  $I_{\mathbf{x}_1}$ 
   $t = \text{torch.rand}(\text{batchsize})$  # Randomly sample  $t \in [0, 1]$ 
   $\mathbf{x}(t) = (1 - t) \times \mathbf{x}_0 + t \times \mathbf{x}_1$  # Sample state at  $t$  based on the linear interpolation
   $\mathbf{x}(t - \tau) = (1 - t + \tau) \times \mathbf{x}_0 + (t - \tau) \times \mathbf{x}_1$  for  $t \in [\tau, 1]$  # Sample delay state at  $t - \tau$  based on the linear interpolation
   $\mathbf{x}(t - \tau) = \mathbf{x}_0 + C_{I_{\mathbf{x}_0} I_{\mathbf{x}_1}} \times (t - \tau)$  for  $t \in [0, \tau]$  # Sample delay state at  $t - \tau$  based on the diverse initial functions
  Loss = {  $v[t, \mathbf{x}(t), \mathbf{x}(t - \tau); \theta] - (\mathbf{x}_1 - \mathbf{x}_0)$  }.pow(2).mean()
  Loss.backward()
  Optimizer.step()
return Model

```

transport trajectories using cubic spline (CSpline) interpolation. These trajectories provide the states, delayed states, and vector fields at various time steps, enabling the computation of the training objective Eq. (16), which is then optimized via backpropagation. Further details on the experimental setup can be found in Table 5.

Details for testing. After training, we evaluate the performance of KP-DFM(C) and KP-CFM on both interpolation and extrapolation tasks. Specifically, we randomly select 100 initial points from the initial distribution $\mathcal{U}(0.2, 1.2)$ and generate predicted trajectories by performing forward integration on the learned vector field, obtaining the predicted trajectories and snapshots at various time steps. For interpolation, we compute the mean 2-Wasserstein distance W_2 between the predicted and ground-truth distributions at $t = 0.5, 1.5, 2.5, \dots, 6.5$. Additionally, since the true dynamics are known, we calculate the error between the predicted and ground-truth trajectories at these points, quantified by the Mean Squared Error (MSE) and Mean Relative Error (MRE). For extrapolation, we similarly compute the W_2 distance between the predicted and ground-truth distributions at $t = 8, 9, 10$, along with the trajectory-wise error using MSE and MRE.

Results. As shown in Table 7 and 8, KP-DFM(C) significantly outperforms CFM in both interpolation and extrapolation tasks, achieving much lower prediction errors for both distributions and individual trajectories, demonstrating superior representational capability. Furthermore, as illustrated in Table 6, CFM fails to reconstruct the underlying dynamics, whereas KP-DFM(C) successfully captures the damped oscillatory dynamics, even when the time delay τ deviates from the ground truth.

Table 6: Additional results on biological autoregulation motif based on CFM with keypoint-guided coupling (KP-CFM), as well as DFM with constant initial functions and keypoint-guided coupling (KP-DFM(C)). For DFM, the generation results are illustrated with various time delays, specifically $\tau = 0.6, 0.8, 1.0$ (ground truth), 1.2, 1.4.

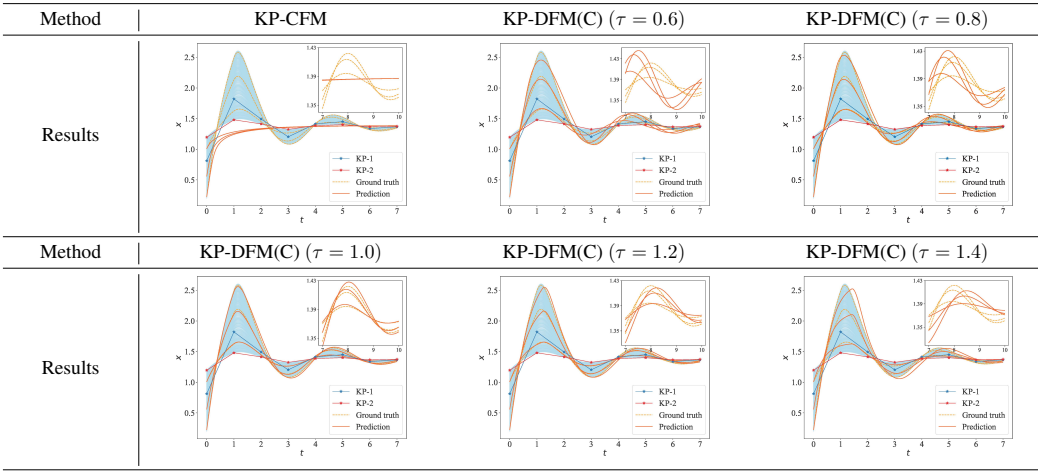


Table 7: Comparison of the interpolation prediction error between KP-CFM and KP-DFM(C) with different time delay τ on the biological autoregulation dataset.

τ	0(CFM)	0.6	0.8	1.0	1.2	1.4
W_2	0.1842	2.82e-2	1.55e-2	1.53e-2	2.27e-2	4.37e-2
MSE	9.21e-2	1.51e-3	4.42e-4	3.45e-4	5.36e-4	2.68e-3
MRE	0.117	2.22e-2	1.23e-2	9.47e-3	1.14e-2	2.57e-2

D.2 SPIRAL DDE

Dataset generation. We consider the following 2-d DDE:

$$\dot{\mathbf{x}}(t) = \mathbf{A} \tanh[\mathbf{x}(t) + \mathbf{x}(t - \tau)], \tag{53}$$

where $\tau = 0.5$ and $\mathbf{A} = [[-1, 20], [-30, -1]] \in \mathbb{R}^{2 \times 2}$. As shown in Fig. 3, the dynamics exhibit trajectory crossings. We sample 1,000 initial points from the uniform distribution $\mathcal{U}(0.8, 1) \times \mathcal{U}(0, 0.1)$ and perform forward integration based on constant initial functions $\psi(h; \mathbf{x}_0) = \mathbf{x}_0, h \in [-0.5, 0]$, to obtain trajectories. After that, snapshots at intervals of 0.05 within $t \in [0, 0.6]$ are selected as training dataset.

Details for training. In each training iteration, 256 data points are randomly and independently sampled from each snapshot. Pairing between the minibatch data points at adjacent time steps is performed using OT. After matching data points across time steps, 256 transport trajectories are

Table 8: Comparison of the extrapolation prediction error between KP-CFM and KP-DFM with different time delay τ on the biological autoregulation dataset.

τ	0(CFM)	0.6	0.8	1.0	1.2	1.4
W_2	2.13e-2	2.25e-2	8.67e-3	3.42e-3	4.55e-3	1.23e-2
MSE	5.51e-4	6.31e-4	1.28e-4	1.46e-5	4.38e-5	2.13e-4
MRE	1.59e-2	1.68e-2	6.75e-3	2.26e-3	3.72e-3	9.37e-3

constructed using CSpline interpolation. These trajectories provide the states, delayed states, and vector fields at various time steps, enabling the computation of the training objective in Eq. (16), which is subsequently optimized via backpropagation. Additional details on the experimental setup are provided in Table 5.

Details for testing. After training, we assess the performance of OT-DFM(C) and OT-CFM on the interpolation task. Specifically, 100 initial points are randomly sampled from the initial distribution $\mathcal{U}(0.8, 1) \times \mathcal{U}(0, 0.1)$, and predicted trajectories are generated via forward integration of the learned vector field, yielding trajectories and snapshots at various time steps. The mean W_2 distance between the predicted and ground-truth distributions is computed at intervals of 0.01 within $t \in [0, 0.6]$. Moreover, we calculate the error between the predicted and ground-truth trajectories at these points, quantified by MSE and MRE.

Results. As shown in Table 10, OT-DFM(C) significantly outperforms CFM in the interpolation prediction task, achieving substantially lower errors for prediction of both distributions and individual trajectories. Additionally, as illustrated in Table 9, CFM collapses around the trajectory crossing area, while OT-DFM(C) successfully captures the spiral dynamics, even when the time delay τ deviates from the ground truth.

Table 9: Additional results on Spiral DDE based on CFM with OT coupling (OT-CFM), as well as DFM with constant initial functions and OT coupling (OT-DFM(C)). For DFM, the generation results are illustrated with various time delays, specifically $\tau = 0.35, 0.40, 0.45, 0.50$ (ground truth), 0.55 .

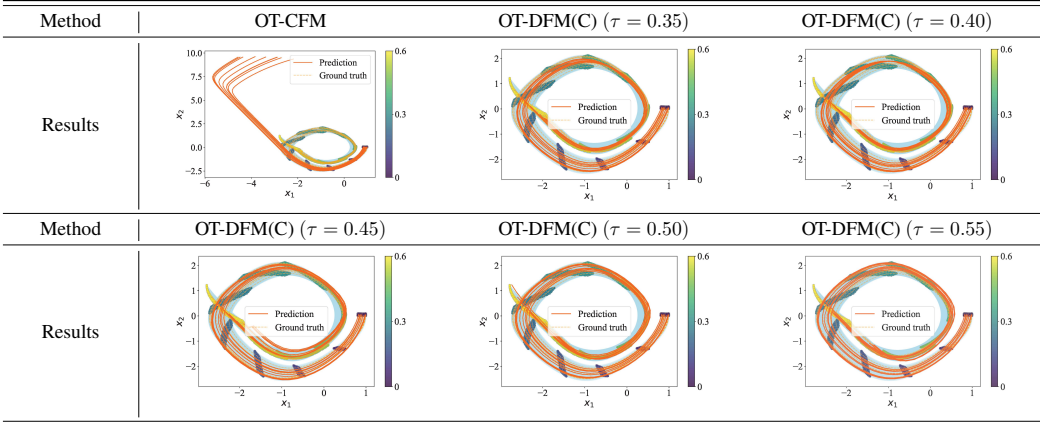


Table 10: Comparison of the interpolation prediction error between OT-CFM and OT-DFM with different time delay τ on the spiral DDE dataset.

τ	0(CFM)	0.35	0.40	0.45	0.50	0.55
W_2	3.95	0.33	0.43	0.42	0.20	0.45
MSE	15.81	0.07	0.13	0.12	0.03	0.14
MRE	13.14	1.25	1.66	1.80	0.68	1.96

D.3 MOUSE HEMATOPOIESIS DATASET

Data preprocessing. To efficiently apply various trajectory inference methods to scRNA-seq data from mouse hematopoiesis, we first project the data into a low-dimensional space. Specifically, we use the reduced two force-directed layouts (SPRING) space for this dataset with batch correction, following the same preprocessing pipeline as TIGON (Sha et al., 2024). Besides, we use the DBSCAN clustering algorithm to categorize all cells at Day 6 into two distinct clusters, corresponding to the Neu fate and Mo fate, respectively.

Selection of initial functions. During training, we match the sampled minibatch cells at Day 2 and 6 using OT. For each matched cell pairs $(\mathbf{x}_0, \mathbf{x}_1)$, where \mathbf{x}_0 is sampled from data of Day 2 and \mathbf{x}_1 is sampled from data of Day 6, the initial function is defined as:

$$\frac{d\psi^*(t; \mathbf{x}_0)}{dt} = \begin{cases} [1, 0], & \text{if } \mathbf{x}_1 \text{ is in Neu cluster,} \\ [0, 1], & \text{if } \mathbf{x}_1 \text{ is in Mo cluster.} \end{cases} \quad (54)$$

Details of training. During training, we align Day 2 and Day 6 with $t = 0$ and $t = 1$, respectively. For testing, the distribution of Day 4 is predicted at $t = 0.5$. In this task, the time delay is set to $\tau = 1$. In each training iteration, we randomly and independently sample 128 cells from each snapshot at Day 2 and 6. Pairing between minibatch data points at adjacent time steps is performed using OT. After matching data points across time steps, 128 transport trajectories γ are constructed using geodesic interpolation (refer to the next paragraph). These trajectories provide the states, delayed states, and vector fields at various time steps, enabling the computation of the training objective in Eq. (16), which is subsequently optimized via backpropagation. Additional details on the experimental setup are provided in Table 5.

Construct geodesic interpolation γ . For each matched cell pairs, geodesic interpolation is employed to generate the interpolation trajectory (Kapuśniak et al., 2024). Here, we first build a data-induced Riemannian metric g on the data manifold \mathcal{M} . Formally, a Riemannian metric g on \mathcal{M} is a smooth family of inner products on the tangent spaces of \mathcal{M} . Specifically, g provides each $\mathbf{x} \in \mathcal{M}$ a positive defined symmetric bilinear form on $T_{\mathbf{x}}\mathcal{M}$,

$$g_{\mathbf{x}} : T_{\mathbf{x}}\mathcal{M} \times T_{\mathbf{x}}\mathcal{M} \rightarrow \mathbb{R}, \quad (55)$$

which induces a norm $\|\cdot\|_{g_{\mathbf{x}}} : T_{\mathbf{x}}\mathcal{M} \rightarrow \mathbb{R}$ defined by $\|v\|_{g_{\mathbf{x}}} = \sqrt{g_{\mathbf{x}}(v, v)}$. Based on this norm, we can calculate the length of any path connecting two points on the manifold and construct the geodesic between them. We aim to equip the ambient space \mathbb{R}^d with an appropriate metric that constrains geodesics to remain close to the data manifold. Specifically, given coordinates, the Riemannian metric g can be equivalently represented as a state-dependent $d \times d$ symmetric positive definite matrix $\mathbf{G}(\mathbf{x})$, such that $g_{\mathbf{x}}(u, v) = u^\top \mathbf{G}(\mathbf{x})v$ for any $u, v \in T_{\mathbf{x}}\mathcal{M}$. Intuitively, to impose manifold constraints on the transport path, $\|\mathbf{G}(\mathbf{x})\|$ is supposed to take smaller values around data points and larger values in regions far from all points in the dataset \mathcal{D} . Therefore, we employ the LAND metric, where $\mathbf{G}(\mathbf{x}; \mathcal{D}) = (\text{diag}(\mathbf{h}(\mathbf{x}; \mathcal{D}) + \epsilon \mathbf{I}))^{-1}$. Here, the components of $\mathbf{h}(\mathbf{x}; \mathcal{D})$ are defined as

$$h_k(\mathbf{x}; \mathcal{D}) = \sum_{i=1}^N (x_k^i - x_k)^2 \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{x} - \mathbf{x}^i\|_2^2\right), \quad (56)$$

where $1 \leq k \leq d$ denotes the k -th component, the superscript i represents the i -th data point in the training dataset \mathcal{D} and σ is the kernel bandwidth. Subsequently, we utilize a neural network ϕ_η to parameterize the geodesic interpolation path between \mathbf{x}_0 and \mathbf{x}_1 as

$$\gamma_\eta(t, \mathbf{x}_0, \mathbf{x}_1) = (1-t)\mathbf{x}_0 + t\mathbf{x}_1 + t(1-t)\phi_\eta(t, \mathbf{x}_0, \mathbf{x}_1). \quad (57)$$

The trajectory γ_η approximates the geodesic between any two points sampled at adjacent time steps on the Riemannian manifold (\mathcal{M}, g) implies learning parameters η to minimize the following loss function:

$$\mathcal{L}_{\mathbf{G}}(\eta) := \mathbb{E}_{\pi(\mathbf{x}_0, \mathbf{x}_1)} \int_0^1 \dot{\gamma}_\eta(t, \mathbf{x}_0, \mathbf{x}_1)^\top \mathbf{G}(\gamma_{\eta, t}; \mathcal{D}) \dot{\gamma}_\eta(t, \mathbf{x}_0, \mathbf{x}_1) dt, \quad (58)$$

where π represents the OT coupling. The optimal parameter is obtained as:

$$\eta^* = \arg \min_{\eta} \mathcal{L}_{\mathbf{G}}(\eta), \quad (59)$$

By substituting η^* into γ_η in Eq. (57), the resulting trajectory γ_{η^*} provides an approximation of the geodesic interpolation on the data manifold.

D.4 iPSCS DATASET

Data Preprocessing. To apply trajectory inference methods to the iPSC scRNA-seq data efficiently, we first reduce the data to a 4-dimensional space using Principal Component Analysis (PCA),

following the preprocessing pipeline of TIGON (Sha et al., 2024). For cells after differentiation (from Day 3), we categorize them into two groups corresponding to the M and En fates using a Gaussian Mixture Model (GMM) with two components. For Day 2, we model the distribution with a single-component GMM and denote its mean as μ_0 . As described above, for Day 5, the distribution is modeled with a two-component GMM, where the means of the M and En fates are denoted as μ_{11} and μ_{12} , respectively.

Selection of initial functions. During training, we match the sampled minibatch cells at Day 2, 4 and 5 using OT. For each matched triplet (x_0, x_1, x_2) , where x_0 is sampled from the Day 2 dataset, x_1 from the Day 4 dataset, and x_2 from the Day 5 dataset, the initial function is defined as:

$$\frac{d\psi^*(t; \mathbf{x}_0)}{dt} = \begin{cases} \mu_{11} - \mu_0, & \text{if } x_2 \text{ is in M component,} \\ \mu_{12} - \mu_0, & \text{if } x_2 \text{ is in En component.} \end{cases} \quad (60)$$

Details of training. During training, we align Day 2, 4 and 5 with $t = 0$, $t = 2$ and $t = 3$, respectively. For testing, the distribution of Day 3 is predicted at $t = 1$. In this task, the time delay is set to $\tau = 3$. In each training iteration, we randomly and independently sample 128 cells from each snapshot at Day 2, 4, and 5. Pairing between minibatch data points at adjacent time steps is performed using OT. After matching data points across time steps, 128 transport trajectories γ are constructed using CSpline interpolation. These trajectories provide the states, delayed states, and vector fields at various time steps, enabling the computation of the training objective in Eq. (16), which is subsequently optimized via backpropagation. Additional details on the experimental setup are provided in Table 5.

D.5 MNIST DATASET

Dataset Generation. We propose a *Semi-paired Image-to-Image Translation* task using the MNIST dataset. The source domain consists of the original MNIST images, while the target domain comprises their corresponding negative images. Specifically, the pixel values X of each image in the original MNIST dataset are inverted by replacing them with $1 - X$, which flips the brightness. The objective is to map each image to its corresponding negative counterpart, as demonstrated in Table 11.

Minibatch coupling. To provide partial supervision, 10% of the training data are paired with their negative counterparts as keypoints. During training, minibatches are independently sampled from the source and target distributions, and KPG-OT coupling (KP-) is applied. Specifically, for each batch, we sample the same number of images from the source distribution and the target distribution. Out of these, around 10% image pairs are selected from the keypoint set, while the remaining images are independently sampled from the respective distributions. It is important to note that the negative counterparts of the remaining images sampled from the initial distribution are not guaranteed to appear in the target batch. Subsequently, we perform KPG-OT pairing on the remaining images in the batch based on the selected keypoint pairs, which are then used for further training.

D.6 CIFAR-10 DATASET

Source distribution. The source distribution for x_0 is defined as follows: with a probability of 50%, x_0 is sampled as $x_0 = \text{torch.randn}(3, 32, 32)/4 - 0.5$, and with the remaining 50% probability, it is sampled as $x_0 = \text{torch.randn}(3, 32, 32)/4 + 0.5$, following the setting in Zhu & Lin (2024). A sample drawn from the source distribution is shown in Table 12.

Trainable initial functions. Designing appropriate initial values for high-dimensional images in image generation tasks is a particularly challenging problem. To address this, we employ `torch.nn.Embedding` to automatically learn the initial functions. Specifically, the source distribution consists of two modules, corresponding to two Gaussian distributions, while the target distribution consists of ten modules, corresponding to ten image categories. Therefore, for each pair of modules (m, n) , where $m \in \{0, 1\}$ and $n \in \{0, 1, \dots, 9\}$, we need to design the time derivative C_{mn} of the initial function (Eq. (22)). To achieve this, we begin by defining `Emb=torch.nn.Embedding(20, 3 × 32 × 32)`, which maps 20 discrete indices to corresponding tensors of the same shape as the image. Then, for each pair (m, n) , we define the initial function as:

$$\frac{d\psi_{mn}^*(t; \mathbf{x}_0)}{dt} = \text{Emb}(10m + n), \quad \psi_{mn}^*(0; \mathbf{x}_0) = \mathbf{x}_0, \quad (61)$$

Table 11: Samples of generation process and results on MNIST data based on CFM with keypoint-guided coupling (KP-CFM), as well as DFM with constant initial functions and keypoint-guided coupling (KP-DFM(C)). For DFM, the generation results are illustrated with various time delays, specifically $\tau = 0.125, 0.25, 0.5,$ and 1.0 . In each image, the first and last columns represent samples obtained from the source data and their negative counterparts, respectively, while the 10 intermediate columns depict the generation process.

Method	KP-CFM	KP-DFM(C) ($\tau = 0.125$)	KP-DFM(C) ($\tau = 0.25$)
Generation			
Method	KP-CFM	KP-DFM(C) ($\tau = 0.5$)	KP-DFM(C) ($\tau = 1.0$)
Generation			

where $h \in [-\tau, 0]$. This approach allows the derivative of the initial functions across different modules to adaptively adjust during training, thereby enabling more flexible and efficient learning.

Experimental setup details. For the image generation tasks, both CFM and DFM are trained following the configurations outlined in Tong et al. (2023a;b). In particular, we utilize a UNet architecture with the following structures and training parameters:

- channels = 128,
- depth = 2,
- channels multiple = [1, 2, 2, 2],
- heads = 4,
- heads channels = 64,
- attention resolution = 16,
- dropout = 0.1,
- batch size per gpu = 128, gpus = 1,
- Adam optimizer with $\beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 10^{-8}$, and no weight decay,
- learning rate = 2×10^{-4} ,
- gradient clipping with norm = 1.0,
- exponential moving average weights with decay = 0.9999.

Table 12: Samples from the source distribution and generation results based on CFM and DFM(D) using independent coupling (I-) or OT coupling (OT-) for adaptive NFE on transporting the Gaussian mixture model with 2 components to the CIFAR-10 dataset.

Method	I-CFM	OT-CFM	I-DFM(D)	OT-DFM(D)
Source				
Generation				

Table 13: Samples from the source distribution and generation results based on CFM and DFM(D) using independent coupling (I-) or OT coupling (OT-) for adaptive NFE on transporting the Gaussian mixture model with 2 components to the CIFAR-10 dataset.

Method	NFE=10	NFE=20	NFE=30	NFE=40
I-CFM				
OT-CFM				
I-DFM(D)				
OT-DFM(D)				

E TRAINING AND INFERENCE TIME OF DFM

In this section, we compare the training and inference times of DFM and CFM. Through detailed analysis and illustrative examples, we show that the training cost of DFM is only marginally higher

than that of CFM. In terms of inference efficiency, when facing trajectory intersections or distributional heterogeneity, DFM achieves substantially faster inference because it produces transport paths that are effectively “straighter” than those of CFM.

Training time. Similar to CFM, we formulate a regression objective to efficiently train the vector field of DFM (see Eq. (16)). The key difference is that DFM additionally requires specifying the initial functions and sampling delayed states. Note that the initial functions are selected and fixed before training and therefore introduce no additional computational cost. The only extra cost arises from delay sampling: in CFM, given sampled initial and target points, we construct an interpolation path and sample the state x at time t . In DFM, after constructing the same interpolation path, we must sample both the state x at time t and the state x_τ at time $t - \tau$. This additional sampling makes DFM slightly more expensive to train than CFM, though it still remains highly efficient in practice. On the 2-d Gaussian mixture dataset, we record the per-iteration training time of both models. Table 14 reports the mean training time per iteration together with the corresponding standard deviation (over 1,000 iterations). DFM requires a slightly longer training time than CFM, but the difference is very small (DFM: 2.014 ± 0.109 s/iter vs CFM: 1.986 ± 0.118 s/iter).

Inference time. Compared to CFM, DFM possesses a stronger representational capacity, enabling it to more accurately approximate the target transport map and the associated trajectories. As a result, during inference, DFM typically follows a much “straighter” transport path, leading to significantly improved inference efficiency. This advantage becomes particularly pronounced in scenarios involving trajectory intersections or distributional heterogeneity. We illustrate this difference using two examples. As shown in Fig. 5, on the 2-d Gaussian mixture dataset, the keypoint-guided optimal transport induces trajectory intersections along the transport process. DFM can directly approximate the linear transport paths and therefore requires only a very small number of Euler discretization steps to achieve accurate generation. In contrast, CFM inevitably induces substantial trajectory bending near the crossing regions, necessitating much more Euler steps to obtain reasonable results. To quantify the inference efficiency, we employ an ODE solver with adaptive step size to compute the generation trajectories of both CFM and DFM. As reported in Table 14, DFM requires only 0.0091 ± 0.0062 s of average inference time per sample, substantially lower than the 0.0442 ± 0.0153 s required by CFM. The 1-d example in Fig. 6 illustrates the inference behavior of DFM and CFM under distributional heterogeneity. Near the singularity, CFM exhibits pronounced trajectory bending, which forces the model to take more discretization steps during inference. In contrast, DFM produces substantially straighter trajectories, achieving accurate generation with essentially a single step. Consequently, DFM attains higher inference efficiency along with improved accuracy in this setting.

Curvature of the transport paths. To assess the straightness of transport trajectories produced by DFM and CFM, we generate 1,000 transport paths on the 2-d Gaussian mixture dataset using the trained DFM and CFM, and compute the absolute curvature at the midpoint of each trajectory. The results, reported in Table 14, show that DFM yields trajectories with consistently small curvature (0.0102 ± 0.0133), indicating substantially straighter paths, whereas CFM produces trajectories with much larger curvature (4.3303 ± 3.9377), reflecting pronounced bending. This observation further supports the higher inference efficiency of DFM compared with CFM.

Generation quality across different steps. In Table 14, we report generation performance of CFM and DFM(C) on the 2-d Gaussian mixture dataset under different numbers of Euler discretization steps. We evaluate generation quality using the W_2 distance between the generated and target distributions. Moreover, because the keypoint-guided coupling strategy induces an optimal transport map ($X \rightarrow -X$), we additionally compute the mean squared error (MSE) between the generated target points and their ground-truth counterparts. The results show that DFM(C) consistently outperforms CFM and achieves high-quality generation with only a small number of steps. On the CIFAR-10 image generation task, we further verify that DFM achieves accurate generation with small NFE (see Table 13), and that its generation quality at low NFE is significantly higher than that of CFM (see Table 4).

Additional remarks. It is important to emphasize that our DFM framework is not designed solely for image generation, where faster inference and straighter transport paths are clear advantages. DFM is also motivated by a broader range of scientific problems—such as inferring single-cell differentiation trajectories from snapshot data or recovering delay-driven dynamical systems—where overly straight trajectories are not only unhelpful but can even be detrimental. Consider single-cell trajectory inference as an example. A fundamental requirement in this setting is that the inferred trajectories

are supposed to remain on the data manifold, ensuring that the recovered differentiation paths reflect biologically plausible intermediate states. Excessively straight transport trajectories tend to leave the manifold, producing intermediate cell states that lack biological meaning and leading to large inference errors. As illustrated in Fig. 4, although MIOFlow (Fig. 4 (b)) infers straighter trajectories than OT-DFM(D) (Fig. 4 (f)), its paths deviate substantially from the gene-expression manifold, resulting in significantly lower trajectory inference accuracy (see Table 2).

Table 14: Comparison of CFM and DFM(C) on the 2-d Gaussian mixture dataset in terms of training time, inference time, transport-trajectory curvature, and generation quality across different discretization steps.

Method	CFM	DFM
Training time	1.986±0.118 s/iter	2.014±0.109 s/iter
Inference time	0.0442±0.0153 s/sample	0.0091±0.0062 s/sample
Curvature	4.3303±3.9377	0.0102±0.0133
1-Step W_2	0.1561	0.1569
1-Step MSE	0.0820	0.0211
2-Step W_2	0.5107	0.1325
2-Step MSE	0.4030	0.0127
4-Step W_2	0.4531	0.0830
4-Step MSE	0.4123	0.0057
Adaptive-Step W_2	0.2884	0.0494
Adaptive-Step MSE	1.0729	0.0024

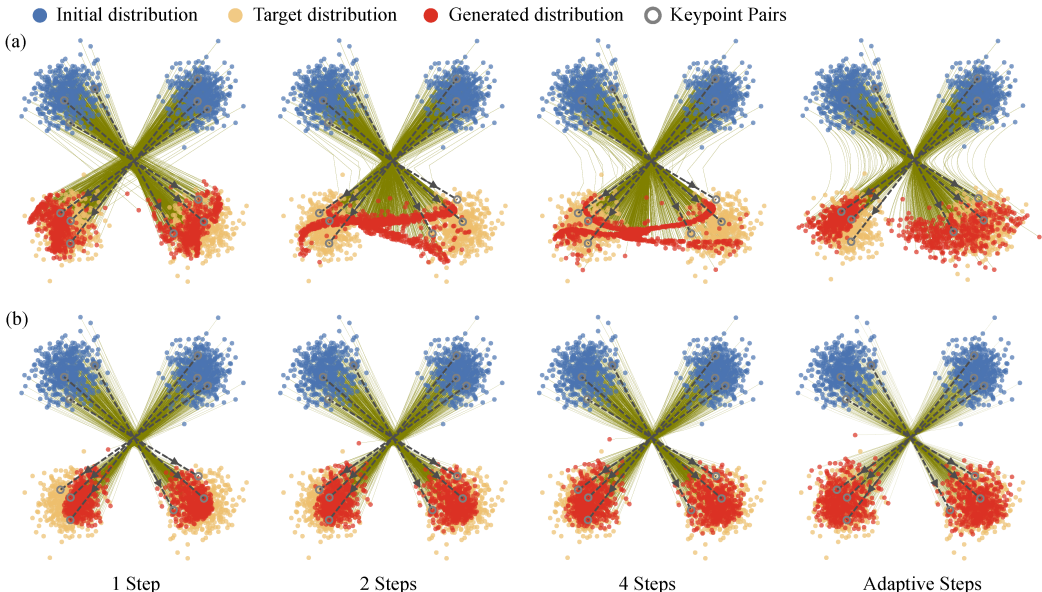


Figure 5: Comparison of CFM (a) and DFM (b) on the Gaussian mixture dataset under different numbers of sampling steps. Since DFM allows trajectory crossings, its generated paths are straighter, enabling accurate generation with very few sampling steps and substantially lower inference time. In contrast, CFM exhibits pronounced bending near trajectory-crossing regions, which requires more sampling steps, leads to slower inference, and results in lower-quality generation.




Figure 6: Comparison of CFM (a) and DFM (b) on the 1-d synthetic dataset with distributional heterogeneity under different numbers of sampling steps. CFM exhibits a singularity near 0, resulting in inaccurate generation and pronounced trajectory bending in the local neighborhood around the singularity. This necessitates more sampling steps and slows down inference. In contrast, DFM, by employing diverse initial functions, produces straighter trajectories for the two subsets from the source to the target distribution, achieving accurate generation with very few sampling steps.

F ABLATION STUDY

In this section, we conduct a systematic ablation study. Recall that, relative to the FM, our DFM framework additionally incorporates the delay term and the initial functions. We conduct sensitivity analyses and discuss selection strategies for both components to demonstrate the robustness of our method to their choice. Moreover, our framework admits different path interpolation schemes, coupling strategies, and the clustering procedure used to determine the initial functions, for which we also provide a detailed examination.

F.1 SENSITIVITY ANALYSIS AND SELECTION STRATEGY FOR THE TIME DELAY τ

To evaluate the effect of time delay τ selection on DFM performance, we perform a sensitivity analysis using both synthetic and real-world datasets.

First, we examine how different values of τ impact the reconstruction of delay dynamical systems. For the biological autoregulation motif, where the true time delay is $\tau = 1$, we conduct extensive experiments with τ values ranging from 0.6 to 1.4. As shown in Table 6, 7 and 8, while the prediction errors for both the distribution and trajectories increase as the delay deviates from the true value, KP-DFM(C) consistently captures the damped oscillation pattern across a wide range of delay values. In contrast, CFM fails to capture this behavior accurately. For the spiral DDE, with a true delay of $\tau = 0.5$, we experiment with delays from 0.35 to 0.55. Although the prediction errors become larger when the delay does not match the true value, OT-DFM(C) consistently outperforms CFM, achieving significantly better results, as shown in Table 10. Additionally, as illustrated in Table 9, OT-DFM(C) successfully captures the spiral dynamics across different delay values, whereas CFM exhibits divergence at the intersections of the trajectories.

Furthermore, we explore the effect of different time delay values on image generation performance. Using the MNIST dataset, we compare the generated images for $\tau = 0.125, 0.25, 0.5$, and 1.0. As shown in Table 11, the generated images from DFM significantly outperform those generated by CFM, with FID scores also notably lower than CFM’s, as summarized in Table 3. Upon examining the impact of varying time delays, we find that performance is least favorable at $\tau = 0.125$, which can be attributed to DFM’s behavior approaching that of CFM as τ tends towards 0. However, for values of $\tau > 0.25$, the generation quality improves substantially.

For the choice of the delay term, we consider two scenarios: (1) the underlying system has no inherent delay; (2) the underlying system is itself a delay dynamical system.

If the system has no explicit delay structure—for example, in image generation tasks—the mere introduction of a delay term is sufficient to mitigate issues such as trajectory intersections, thereby improving generation quality and yielding results that significantly outperform ODE-based FM models. The sensitivity analysis of the delay term on the MNIST dataset further supports this claim

(see Table 3). In such cases, the key factor is simply the *presence* of a delay term, whereas the exact numerical value of τ plays a relatively minor role. If an optimal value of τ is desired, standard hyperparameter selection techniques such as grid search can be employed.

If the underlying data correspond to a genuinely delay system, then the choice of τ becomes critical for accurately capturing the system’s dynamics, as illustrated in the biological autoregulation example (see Table 6, 7 and 8). Here, the closer τ is to its true value, the more accurate the trajectory reconstruction. For certain biophysical systems, prior domain knowledge often provides approximate estimates of τ (an der Heiden, 1979; Glass et al., 2021; Rihan et al., 2021). When such prior information is unavailable, we can resort to delay–identification methods from dynamical systems theory to infer the delay structure directly from data (Ma et al., 2010; 2017).

F.2 SENSITIVITY ANALYSIS AND SELECTION STRATEGY FOR THE INITIAL FUNCTIONS

For generative frameworks based on DDEs, the initial function is an essential component. In the main text, we construct different forms of initial functions for different tasks, such as constant initial functions and diverse initial functions. For the diverse initial function, we further specify different constant time derivatives depending on the scenario (see Section 4.3.2, Eq. (22)). In this section, we demonstrate that DFM exhibits strong robustness with respect to both the choice of initial-function form and the selection of constant derivatives.

We first examine how to select the form of the initial functions. Note that diverse initial-function designs are introduced to better address transport across heterogeneous distributions. Distinct initial functions guide samples toward different target subsets. Therefore, if the source and target distributions are known to decompose into multiple modules (e.g., when single cells differentiate into several cell types), it is preferable to employ the DFM with diverse initial functions to model the generative process. If no distributional heterogeneity is present, a simple constant initial function suffices. In fact, we have theoretically established that the DFM with constant initial functions already possesses universal approximation capability for continuous transport maps between arbitrary distributions. Moreover, experiments on synthetic delay–dynamical data and the MNIST dataset further confirm that a constant initial function enables DFM to model transport maps with trajectory intersections and to capture the underlying delay-driven dynamical behavior.

Next, we examine how the selection of constant derivatives affects the performance of DFM when diverse initial functions are used. We conduct a series of comparative experiments on the real iPSCs single-cell trajectory inference task. In the main text, the initial functions employed correspond to Eq. (60), which are specifically designed based on the data structure. Here, we instead adopt a randomized procedure to generate distinct initial functions for different cell types, as follows:

$$\frac{d\psi^*(t; \mathbf{x}_0)}{dt} = \begin{cases} \text{RandomVec1}, & \text{if } \mathbf{x}_2 \text{ is in M component,} \\ \text{RandomVec2}, & \text{if } \mathbf{x}_2 \text{ is in En component.} \end{cases} \quad (62)$$

where $\text{RandomVec1}=\text{torch.rand}(4)$ and $\text{RandomVec2}=\text{torch.rand}(4)$ are 4-dimensional vectors drawn independently from the uniform distribution. We randomly sampled 10 sets of initial functions, trained a separate DFM for each, and conducted the single-cell trajectory inference task. The mean and standard deviation of the inference errors are reported in Table 15 (corresponding to DFM(DR)). The results show that, compared with the specially designed initial functions (Eq. (60)), using randomized initial functions yields only a slight increase in the inference error, while still outperforming the DFM with constant initial functions and substantially surpassing the performance of CFM. This demonstrates the strong robustness of the DFM framework to the choice of initial functions. Moreover, the comparison between DFM(D) and DFM(C) further indicates that, in the presence of distributional heterogeneity, adopting diverse initial functions leads to improved generative performance, thereby reinforcing the selection strategy discussed above.

F.3 ABLATION STUDY ON THE QUALITY OF CLUSTERING

For DFM with diverse initial functions, the design of these functions depends on the clustering results of the source and target samples. We examine two scenarios in which the clustering results are suboptimal and analyze how this affects model performance: (1) the number of clusters is smaller than the true number of underlying classes, and (2) the number of clusters exceeds the true number of classes.

Table 15: Trajectory inference results on the qPCR iPSC dataset with different initial functions. DFM(DR) denotes the DFM model equipped with diverse initial functions, where the corresponding constant time derivatives of the initial functions are randomly sampled.

Method	$W_2(L)$	MMD(L)	$W_2(F)$	MMD(F)
OT-CFM	0.579	0.492	0.226	0.030
OT-DFM(C)	0.553	0.447	0.234	0.041
OT-DFM(D)	0.532	0.399	0.213	0.027
OT-DFM(DR)	0.545 ± 0.007	0.441 ± 0.036	0.218 ± 0.008	0.028 ± 0.007

When the number of clusters is smaller than the true number of classes, multiple distinct subsets share the same initial function. In this case, the distributional heterogeneity among these subsets cannot be fully resolved. Nonetheless, the presence of the delay term mitigates the resulting singularities, and the overall generation quality remains substantially better than that of CFM. We illustrate this with a 1-d example in which the task is to transport $\mathcal{U}(-1, 1)$ to the two-component mixture $\frac{1}{2}\mathcal{U}(-3, -2) + \frac{1}{2}\mathcal{U}(2, 3)$. When the two modules are correctly identified and assigned distinct initial functions, the generated distribution (Fig. 7 (c)) closely matches the target. However, if the two modules are incorrectly grouped into a single cluster, the resulting generation (Fig. 7 (b)) exhibits the singularity problem near 0 in the source domain, causing a degradation in the generation quality. Nevertheless, the delay term still alleviates this issue, yielding results that remain superior to those of CFM.

When the number of clusters exceeds the true number of classes, a single true class may be split into two subclusters. In this situation, the over-segmentation provides sufficient refinement to resolve the heterogeneity-induced singularities. As shown in Fig. 7 (d), further subdividing one module of the target distribution does not negatively impact the final generation quality.




Figure 7: The generated results of CFM and DFM under different numbers of clusters, with the squared 2-Wasserstein distance between the generated and target distributions.

F.4 SELECTION OF OTHER COMPONENTS IN THE DFM FRAMEWORK

In this section, we outline the rationale behind our choices of path-interpolation methods and coupling strategies. Across different experiments, we adopt different interpolation schemes and coupling strategies, highlighting the flexibility of the DFM framework and its ability to accommodate a variety of modeling choices. Both the selection of interpolation methods and the choice of coupling strategies are determined by the characteristics of each task and align with prior work in related domains.

Specifically, for path interpolation, we use linear interpolation when no prior geometric constraints are present, as it yields the most direct transport path and improves generative efficiency (Liu et al., 2023; Esser et al., 2024). For tasks such as single-cell trajectory inference, where paths are required to remain on the underlying data manifold, we adopt geodesic interpolation, consistent with prior work (Huguet et al., 2022; Kapuśniak et al., 2024). For continuous transport involving three or more distributions, paths connecting multiple intermediate points must be constructed, for which we employ cubic spline interpolation (Pooladian et al., 2024).

For the choice of coupling strategy, independent coupling or optimal-transport coupling is typically used when no prior constraints are available (e.g., in image generation tasks). In contrast, for certain biological systems such as single-cell gene expression, many work has shown that their evolution tends to follow energy-efficient transport patterns (Bunne et al., 2024). In such cases, employing

the optimal-transport coupling strategy is the most appropriate choice and has been widely adopted in single-cell trajectory inference (Kapuśniak et al., 2024; Sha et al., 2024). Moreover, in many scenarios, partial keypoint correspondences are known a priori, and the coupling strategy must respect these constraints. A line of research has proposed keypoint-guided optimal transport (Gu et al., 2022; 2023), which preserves keypoint correspondences while still minimizing the overall transport cost. Therefore, when partial prior coupling information is available, keypoint-guided optimal transport constitutes the most suitable strategy.

G COMPARISON OF DFM WITH OTHER ODE-BASED GENERATIVE METHODS

In Section 5, we discuss a series of works that modify ODE-based generative frameworks to overcome FM’s limitations in handling trajectory intersections and distributional heterogeneity. Here, we focus on two representative methods: Augmented Bridge Matching (AugBM) (De Bortoli et al., 2023) and Switched Flow Matching (SFM) (Zhu & Lin, 2024). AugBM allows trajectory intersections during transport by augmenting the velocity field with information from the initial point. SFM introduces an additional switching signal as a conditioning variable and employs switching ODEs to enable transport between different modules of the source and target distribution, thereby alleviating singularities caused by distributional heterogeneity.

As discussed in Section 5, these ODE-based modifications typically address only one type of limitation and cannot simultaneously resolve both. We illustrate this with two simple examples. In Fig. 8 (a-c), we compare the generation results of CFM, DFM(C), and SFM trained under the keypoint-guided optimal transport strategy on the 2-d Gaussian mixture dataset. While SFM can alleviate issues caused by distributional heterogeneity, its representation capacity is still weaker than DFM and cannot perfectly resolve trajectory intersections. Notably, although SFM allows trajectories to cross between different mixture components, trajectories within the same component still cannot intersect. This is because points sampled from the same component share the same switching signal during generation, resulting in a single ODE that forbids intra-component trajectory crossing. Consequently, SFM cannot precisely realize the optimal transport mapping for this task ($X \rightarrow -X$). In Fig. 8 (d-f), we compare CFM, DFM(D), and AugBM on the 1-d example with distributional heterogeneity. Although AugBM allows trajectory intersections during transport, it still suffers from singularities when transporting between heterogeneous distributions, preventing it from achieving precise distribution transport.

We further highlight the distinctions between these ODE-based extensions and our proposed DFM framework. The ODE-based methods remain confined to the ODE generative paradigm, typically addressing their limitations either by augmenting the vector field with additional conditioning variables or by designing highly specialized transport paths and procedures. In contrast, DFM departs from the ODE framework and leverages the vector field with a delay term to realize distributional transport. Notably, the inclusion of the delay term naturally overcomes the limitations inherent in ODE-based approaches within the original phase space, without requiring auxiliary conditioning variables or carefully engineered transport trajectories. Moreover, for transport tasks whose underlying dynamics involve intrinsic delays, such as the gene autoregulation system in Section 6.2, ODE-based generative frameworks are incapable of recovering the true dynamics, whereas DFM, based on DDEs, can accurately model such delayed dynamics.

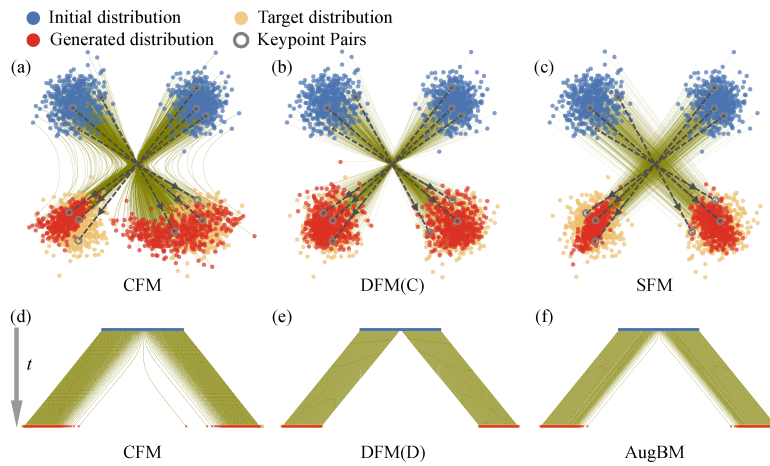


Figure 8: (a–c) Comparison of generation results from CFM (a), DFM(C) (b), and SFM (c) on the 2-d Gaussian mixture dataset. (d–f) Comparison of generation results from CFM(d), DFM(D) (e), and AugBM (f) on the 1-d dataset with distributional heterogeneity.