Unlocking the Duality between Flow and Field Matching

Anonymous Authors[1]

Abstract

Conditional Flow Matching (CFM) unifies conventional generative paradigms such as diffusion and flow-based models. Interaction Field Matching (IFM) is a recently proposed framework that generalizes Electrostatic Field Matching (EFM), rooted in Poisson Flow Generative Models (PFGM). While both frameworks define generative dynamics, they start from different objects: CFM specifies a conditional probability path in data space, whereas IFM specifies a physics-inspired interaction field in an augmented data space. This raises a basic question: **are CFM and IFM genuinely different, or are they two descriptions of the same underlying dynamics?** We show that they coincide for a natural subclass of IFM, which we call *forward-only* IFM. Specifically, we construct mappings between CFM and forward-only IFM and show that they induce the same generative dynamics. We further show that general IFM is strictly more expressive: it includes EFM and other interaction fields that cannot be realized within the standard CFM formulation. Our findings suggest developing generative models using both interpretations rather than treating them separately. Moreover, they highlight a novel direction for generative modeling based on backward-oriented field lines, which lies outside the conventional CFM formalism and may lead to new generative properties.

1 Introduction

Modern continuous-time generative modeling is largely organized around flow-based descriptions of transport. Diffusion (Sohl-Dickstein et al., 2015; Song et al., 2020) and flow-based (Liu et al., 2022; Lipman et al., 2022; Albergo & Vanden-Eijnden, 2022) generative models can both be unified under the Conditional Flow Matching (Tong et al., 2023, CFM) framework. CFM constructs a generative ODE by specifying a conditional probability path in data space together with a compatible conditional velocity field, and then marginalizing these conditional objects to obtain global dynamics that transport the source distribution to the target distribution (Figure 1b).

A separate line of work has developed physics-inspired (Liu et al., 2023; Shlenskii & Korotin, 2026), field-based generative models, including Poisson Flow Generative Models (Xu et al., 2022, PFGM), PFGM++ (Xu et al., 2023), and Electrostatic Field Matching (Kolesov et al., 2025, EFM). Rather than starting from probability paths in data space, these methods define vector fields in an augmented space and generate samples by following their field lines. This perspective is formalized by the recently proposed Interaction Field Matching (Manukhov et al., 2026, IFM) framework, which constructs particle-dependent interaction fields and combines them via a superposition principle to obtain a global transport field (Figure 1c,d).

Thus, CFM and IFM define generative dynamics from different first principles: CFM is path-first and probabilistic, whereas IFM is field-first and geometric. Understanding whether these views are genuinely distinct is important because it determines whether techniques, objectives, and intuitions developed in one framework can be transferred to the other. This raises a basic question: *are CFM and IFM fundamentally distinct, or do they produce the same generative dynamics expressed in different languages?*

Our main contributions are as follows:

- **Flow–field duality.** We identify forward-only IFM as the field-theoretic counterpart of CFM and prove that the two frameworks induce the same ODE-based generative dynamics (§3.2, §3.3).
- **A strict separation between CFM and general IFM.** We show that general IFM is strictly more expressive than CFM (§4), because it includes EFM and other interaction fields with backward-oriented lines that cannot be realized within the standard CFM formulation.

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the FoGen Workshop at ICML 2026. Do not distribute.

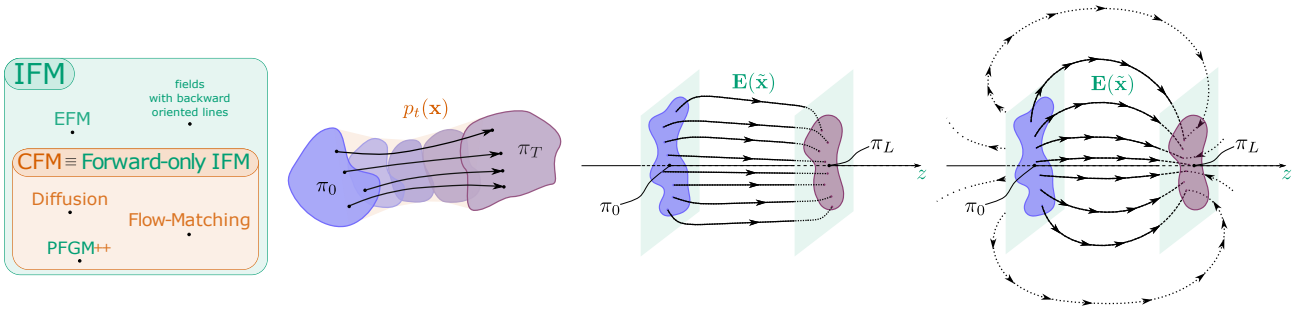

(a) Flow-field taxonomy (b) CFM (c) Forward-only IFM ≡ CFM (d) EFM ∈ IFM \ CFM

Figure 1. **Overview of the flow–field duality.** We establish that CFM and forward-only IFM induce the same generative dynamics (§3), yielding the equivalence CFM ≡ forward-only IFM. Panel (a) summarizes the resulting taxonomy. While conventional CFM (b, §2.1) admits an equivalent forward-only field representation (c, §2.2), general IFM is strictly larger and includes fields with backward-oriented field lines, such as EFM (d, §2.2), which lie outside standard CFM.

- **Explicit bidirectional constructions.** We provide explicit mappings in both directions: from conditional flows to forward-only interaction fields (§3.2), and from forward-only interaction fields back to conditional flows (§3.3). These constructions yield a unified flow–field view of the two frameworks, enabling both probabilistic and geometric interpretations of the same dynamics.

These findings are important for two reasons. First, they suggest developing generative models using the probabilistic and geometric interpretations jointly, rather than treating them as separate frameworks. Second, they identify backward-oriented field lines as a key mechanism by which general IFM goes beyond conventional CFM, opening a new direction for field-based generative modeling outside the standard conditional flow matching formalism.

2 Preliminaries: CFM vs. IFM?

In this section, we review the two constructions of generative dynamics that will be compared throughout the paper. Both frameworks define dynamics of the form $d\mathbf{x} = \mathbf{v}_t(\mathbf{x})\,dt$, which transport a source/prior distribution, denoted by π_T or π_L depending on the framework's notation, to a target distribution π_0. First, in §2.1, we recall CFM, a well-established framework that generalizes diffusion and flow-based generative models (Gao et al., 2025). Next, in §2.2, we review the recently proposed IFM, which can be viewed as a generalization of electrostatic generative models. In §2.3, we provide a side-by-side comparison of the key ingredients of the two approaches. Finally, in §2.4, we discuss PFGM, which, to our knowledge, is the only existing framework that has been found to admit both flow and field descriptions. Throughout this section, we use **orange** and **green** to highlight terms specific to CFM and IFM, respectively.

2.1 Conditional Flow Matching (CFM)

Intuition. The CFM framework begins by specifying two conditional objects in the data space $\mathbb{R}^D$: **(a)** a **conditional distribution path** $p_t^{\mathbf{x}_0,\mathbf{x}_T}$ that interpolates between $\mathbf{x}_0 \in \mathbb{R}^D$ and $\mathbf{x}_T \in \mathbb{R}^D$, i.e., $p_{t=0}^{\mathbf{x}_0,\mathbf{x}_T} = \delta(\mathbf{x}_0)$ and $p_{t=T}^{\mathbf{x}_0,\mathbf{x}_T} = \delta(\mathbf{x}_T)$; and **(b)** a compatible **conditional velocity** $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) : \mathbb{R}^D \to \mathbb{R}^D$ associated with this path. In the two-sided setting, both objects are conditioned on a target–source sample pair $\mathbf{x}_0, \mathbf{x}_T$. Marginalizing these conditional objects with respect to a given coupling $\pi_{0,T}(\mathbf{x}_0, \mathbf{x}_T)$ yields a global velocity $\mathbf{v}_t$, which induces the generative dynamics $d\mathbf{x} = \mathbf{v}_t(\mathbf{x})\,dt$, used to transport samples from the source distribution π_T to the target distribution π_0.

We use blue to emphasize that the elementary CFM objects $p_t^{\mathbf{x}_0,\mathbf{x}_T}$ and $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}$ can be constructed in either a *one-sided* or *two-sided* manner. In the one-sided case, $p_t^{\mathbf{x}_0}$ and $\mathbf{v}_t^{\mathbf{x}_0}$ are defined using only target samples $\mathbf{x}_0 \sim \pi_0$. In the two-sided case, $p_t^{\mathbf{x}_0,\mathbf{x}_T}$ and $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}$ are defined from paired target and source samples $\mathbf{x}_0, \mathbf{x}_T \sim \pi_{0,T}$, for example under the independent coupling $\pi_{0,T} = \pi_0 \times \pi_T$. This choice does not affect the theoretical guarantees (Tong et al., 2023). The one-sided setting is common in noise-to-data generation (Song et al., 2020), where the source distribution π_T is a simple prior, such as a standard Gaussian, independent of the target distribution π_0. By contrast, the two-sided setting is more typical for data-to-data translation (Liu et al., 2022; Albergo et al., 2023b).

Definition (Conditional Flow). A *Conditional Flow* (Tong et al., 2023) is specified by a pair $(\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}), p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}))$ consisting of a conditional velocity field and a path of conditional distributions. Both objects depend on $\mathbf{x}_0, \mathbf{x}_T$ and satisfy the ***continuity equation*** $\forall \mathbf{x} \in \mathbb{R}^D$ and $t \in (0, T)$

$$\partial_t p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) + \nabla_{\mathbf{x}} \cdot \left(\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\right) = 0, \quad (1)$$

where $\nabla_{\mathbf{x}} \cdot$ denotes the divergence operator with respect to

$\mathbf{x}$, together with the *boundary conditions*

$$\lim_{t\to 0} p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \delta_{\mathbf{x}_0}(\mathbf{x}),\ \lim_{t\to T} p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \delta_{\mathbf{x}_T}(\mathbf{x}), \quad (2)$$

that hold in the sense of weak convergence.

Generative dynamics. This pointwise-defined conditional flow induces a generative model that transports a source distribution $\pi_T(\mathbf{x})$ to a target distribution $\pi_0(\mathbf{x})$. First, marginalizing the conditional path with respect to the data distribution $\pi_{0,T}$ yields a **path of marginal distributions**

$$p_t(\mathbf{x}) = \int p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, \pi_{0,T}(\mathbf{x}_0,\mathbf{x}_T)\, d\mathbf{x}_0\, d\mathbf{x}_T, \quad (3)$$

which satisfies $p_{t=0} = \pi_0$ and $p_{t=T} = \pi_T$. Second, one can show that the marginals p_t are generated by a deterministic flow governed by the ODE

$$d\mathbf{x}/dt = \mathbf{v}_t(\mathbf{x}) := \mathbb{E}_{p_t(\mathbf{x}_0,\mathbf{x}_T|\mathbf{x}_t=\mathbf{x})} \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}), \quad (4)$$

where the **global velocity** $\mathbf{v}_t(\mathbf{x})$ is the conditional expectation of the conditional velocity $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$ over endpoint pairs $\mathbf{x}_0, \mathbf{x}_T$, given the intermediate state $\mathbf{x}_t = \mathbf{x}$.

Although the drift expression (4) is explicit, it is generally intractable to evaluate in practice, since sampling from the conditional distribution $p_t(\mathbf{x}_0, \mathbf{x}_T \mid \mathbf{x}_t)$ is infeasible.

Training. The conditional expectation in (4) naturally yields a regression objective for estimating the drift. In particular, $\mathbf{v}_t(\mathbf{x})$ can be learned by minimizing

$$\mathbb{E}_{t,\mathbf{x}_0,\mathbf{x}_T,\mathbf{x}_t} \left\| \mathbf{v}_t(\mathbf{x}_t) - \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}_t) \right\|_2^2, \quad (5)$$

where $t \sim \mathcal{U}[0,T]$, $\mathbf{x}_t \sim p_t^{\mathbf{x}_0,\mathbf{x}_T}$, and the conditional $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}$ provides a **single-sample target**.

Instantiation I (Two-sided Stochastic Interpolant). One of the most common instances of the CFM is the two-sided Stochastic Interpolant (Albergo et al., 2023a, §2.1), defined by $\mathbf{x}_t = \mathcal{I}(t,\mathbf{x}_0,\mathbf{x}_T) + \sigma(t)\,\epsilon$, where **(a)** $\mathcal{I}(0,\mathbf{x}_0,\mathbf{x}_T) = \mathbf{x}_0$ and $\mathcal{I}(T,\mathbf{x}_0,\mathbf{x}_T) = \mathbf{x}_T$; **(b)** $\sigma(0) = 0$, $\sigma(T) = 0$, and $\sigma(t) > 0$ for all $t \in (0,T)$; **(c)** $\epsilon \sim \mathcal{N}(0,I)$ is independent of $\mathbf{x}_0, \mathbf{x}_T$. This construction induces a path of conditional distributions $p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \mathcal{N}\big(\mathcal{I}(t,\mathbf{x}_0,\mathbf{x}_T), \sigma(t)^2 I\big)$, with the associated conditional velocity $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}_t) = \dot{\mathcal{I}}(t,\mathbf{x}_0,\mathbf{x}_T) + (\dot\sigma(t)/\sigma(t))\big(\mathbf{x}_t - \mathcal{I}(t,\mathbf{x}_0,\mathbf{x}_T)\big)$, where a dot over a function denotes the partial derivative with respect to time.

Instantiation II (One-sided Stochastic Interpolant). Another commonly used instance is the one-sided Stochastic Interpolant (Albergo et al., 2023a, §3.2), also known as Flow Matching (Lipman et al., 2022), defined by $\mathbf{x}_t = J(t,\mathbf{x}_0) + \sigma(t)\,\epsilon$, where **(a)** $J(0,\mathbf{x}_0) = \mathbf{x}_0$ and $J(T,\mathbf{x}_0) = \mathbf{0}$; **(b)** $\sigma(0) = 0$, $\sigma(T) = 1$, and $\sigma(t) > 0$ for all $t \in (0,T)$; **(c)** $\mathbf{x}_T \sim \mathcal{N}(0,I)$ is independent of $\mathbf{x}_0$. As in the two-sided case, this yields a path of conditional distributions $p_t^{\mathbf{x}_0}(\mathbf{x}) = \mathcal{N}\big(J(t,\mathbf{x}_0), \sigma(t)^2 I\big)$, with the corresponding conditional velocity $\mathbf{v}_t^{\mathbf{x}_0}(\mathbf{x}_t) = \dot J(t,\mathbf{x}_0) + (\dot\sigma(t)/\sigma(t))\big(\mathbf{x}_t - J(t,\mathbf{x}_0)\big)$.

Remark. Although this construction conditions only on $\mathbf{x}_0$, it still specifies a right-endpoint prior π_T that is independent of π_0, namely the standard Gaussian: $\pi_T(\mathbf{x}_T \mid \mathbf{x}_0) = \pi_T(\mathbf{x}_T) = \mathcal{N}(0,I)$. Moreover, in the special case $J(t,\mathbf{x}_0) = \mathbf{x}_0(1 - t/T)$ and $\sigma(t) = t/T$, one recovers the classical linear interpolant in (Lipman et al., 2022). The same choice is also closely related to the Gaussian flow-matching formulation of diffusion models (Gao et al., 2025).

2.2 Interaction Field Matching (IFM)

Intuition. Interaction Field Matching (IFM) lifts the data space to an extended space $\mathbb{R}^{D+1}$. Each data point $\mathbf{x} \in \mathbb{R}^D$ is augmented with an additional coordinate $z \in \mathbb{R}$, giving $\tilde{\mathbf{x}} = (\mathbf{x}, z)$. The target and source distributions are placed on two parallel plates: the target distribution π_0 lies on the plate $z = 0$, while the source distribution π_L lies on the plate $z = L > 0$. Thus, samples $\mathbf{x}_0 \sim \pi_0$ and $\mathbf{x}_L \sim \pi_L$ are identified with the augmented points $\tilde{\mathbf{x}}_0 = (\mathbf{x}_0, 0)$ and $\tilde{\mathbf{x}}_L = (\mathbf{x}_L, L)$. In EFM, this geometry is interpreted as a $(D+1)$-dimensional electric capacitor.

The elementary object in IFM is an **interaction field** $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}}) : \mathbb{R}^{D+1} \to \mathbb{R}^{D+1}$, defined for an endpoint pair $\mathbf{x}_0, \mathbf{x}_L = (\mathbf{x}_0, \mathbf{x}_L)$. Its field lines connect the two augmented endpoints $\tilde{\mathbf{x}}_0$ and $\tilde{\mathbf{x}}_L$. A **global field** $\mathbf{E}(\tilde{\mathbf{x}})$ is then obtained by **superposition**: the interaction fields are averaged over endpoint pairs $(\mathbf{x}_0, \mathbf{x}_L) \sim \pi_{0,L}$. The field lines of the resulting global field define the transport from the source distribution π_L to the target distribution π_0.

Remark. In contrast to conditional flows, interaction fields are inherently two-sided: their field lines are required to connect two particles located at $\tilde{\mathbf{x}}_0$ and $\tilde{\mathbf{x}}_L$. This is reflected in our notation: for conditional flows, the source conditioning point $\mathbf{x}_T$ is shown in blue, whereas for interaction fields the source point $\mathbf{x}_L$ is treated on the same basis as the target point $\mathbf{x}_0$ and is therefore shown in black.

Definition (Interaction Field). A vector field $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}}) \in \mathbb{R}^{D+1}$ is an *Interaction Field* (Manukhov et al., 2026) if it satisfies the following three properties.

First, its ***field lines*** *connect the two particles*:

$$\begin{cases} \dfrac{d\tilde{\mathbf{x}}(\tau)}{d\tau} = \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}}(\tau)), \\ \tilde{\mathbf{x}}(\tau_s) = \tilde{\mathbf{x}}_0, \quad \tilde{\mathbf{x}}(\tau_f) = \tilde{\mathbf{x}}_L, \end{cases} \quad (6)$$

where τ_s and τ_f denote the start and end values of the field-line parameter, respectively.

Remark. The parameter τ in (6) is not a global schedule like the time variable in Conditional Flows (1). Rather, τ is an arbitrary scalar parameter used solely to trace a field line (integral curve) of $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})$.

Second, the field is ***divergence-free (flux-conserving)***:

$$\nabla \cdot \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}}) = 0, \quad \forall \tilde{\mathbf{x}} \in \mathbb{R}^{D+1} \backslash \{\tilde{\mathbf{x}}_0, \tilde{\mathbf{x}}_L\} \tag{7}$$

Third, for any closed surface ∂M enclosing $\tilde{\mathbf{x}}_0$ but not $\tilde{\mathbf{x}}_L$, the total flux

$$\int_{\partial M} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}}) \cdot \mathbf{dS} \;=\; \Phi_0 \tag{8}$$

is independent of the particle locations $\tilde{\mathbf{x}}_0, \tilde{\mathbf{x}}_L$, i.e., the field has a ***fixed total flux***.

Special case (forward-only Interaction Field). We say that an interaction field is *forward-only* if **(a)** it is defined in the region between the plates ($0 \le z \le L$), and, **(b)** its z-component satisfies $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})_z > 0$ everywhere the field is nonzero.

Generative dynamics. Given the target and source distributions π_0 and π_L, the induced **global field** $\mathbf{E}(\tilde{\mathbf{x}})$ is defined via a ***generalized superposition principle*** as

$$\mathbf{E}(\tilde{\mathbf{x}}) = \int \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})\, \pi_{0,L}(\mathbf{x}_0, \mathbf{x}_L)\, d\mathbf{x}_0\, d\mathbf{x}_L. \tag{9}$$

That is, $\mathbf{E}(\tilde{\mathbf{x}})$ is the average contribution of interaction fields over particle pairs. As shown in (Manukhov et al., 2026), one can transport samples from π_0 to π_L by following the field lines of $\mathbf{E}$:

$$\frac{d\tilde{\mathbf{x}}(\tau)}{d\tau} = \mathbf{E}(\tilde{\mathbf{x}}(\tau)). \tag{10}$$

In general, simulating (10) poses practical difficulties: **(a)** the boundary conditions are implicit because the corresponding τ_s and τ_f are unknown; **(b)** field lines may move both forward ($\mathbf{E}(\tilde{\mathbf{x}})_z > 0$) and backward ($\mathbf{E}(\tilde{\mathbf{x}})_z < 0$) along the z-axis; and **(c)** field lines may be defined beyond the capacitor region, $0 \le z \le L$. For a forward-only interaction field, the global field is forward-only as well, i.e., $\mathbf{E}_z > 0$, and (10) can be reparameterized using the physically meaningful coordinate z:

$$\frac{d\tilde{\mathbf{x}}}{dz} = \frac{d\tilde{\mathbf{x}}}{d\tau} \cdot \frac{d\tau}{dz} = \mathbf{E}(\tilde{\mathbf{x}})\, \mathbf{E}(\tilde{\mathbf{x}})_z^{-1} = \left(\frac{\mathbf{E}(\tilde{\mathbf{x}})_{\mathbf{x}}}{\mathbf{E}(\tilde{\mathbf{x}})_z}, 1 \right). \tag{11}$$

This form yields explicit boundaries, $z(\tau_s) = 0$ and $z(\tau_f) = L$, and induces generative dynamics

$$\frac{d\mathbf{x}}{dz} = \frac{\mathbf{E}(\tilde{\mathbf{x}})_{\mathbf{x}}}{\mathbf{E}(\tilde{\mathbf{x}})_z}, \tag{12}$$

which is defined in the data space $\mathbb{R}^D$ and z here plays the role of a time-like variable.

Training. For a general Interaction Field, one has to estimate the full field $\mathbf{E}(\tilde{\mathbf{x}}) \in \mathbb{R}^{D+1}$ at every point $\tilde{\mathbf{x}} \in \mathbb{R}^{D+1}$ where the field is defined. For forward-only fields, however, it suffices to estimate the *normalized* field $\mathbf{E}(\tilde{\mathbf{x}})/\|\mathbf{E}(\tilde{\mathbf{x}})\|_2$. Indeed, the generative dynamics (12) can be rewritten as

$$\frac{d\mathbf{x}}{dz} = \frac{\mathbf{E}(\tilde{\mathbf{x}})_{\mathbf{x}}}{\mathbf{E}(\tilde{\mathbf{x}})_z} = \frac{\mathbf{E}(\tilde{\mathbf{x}})_{\mathbf{x}}}{\|\mathbf{E}(\tilde{\mathbf{x}})\|_2} \cdot \frac{\|\mathbf{E}(\tilde{\mathbf{x}})\|_2}{\mathbf{E}(\tilde{\mathbf{x}})_z}. \tag{13}$$

Motivated by this observation, (Manukhov et al., 2026) use the following training objective:

$$\mathbb{E}_{\tilde{\mathbf{x}} \sim p_{\text{vol}}} \left\| f_\theta(\tilde{\mathbf{x}}) - \frac{\mathbf{E}(\tilde{\mathbf{x}})}{\|\mathbf{E}(\tilde{\mathbf{x}})\|_2} \right\|_2^2, \tag{14}$$

where **(a)** $f_\theta : \mathbb{R}^{D+1} \to \mathbb{R}^{D+1}$ is a neural network that estimates the normalized Interaction Field; **(b)** p_{vol} is a heuristic **distribution intended to cover the volume** between the plates at $z = 0$ and $z = L$, where the field is defined; and **(c)** $\mathbf{E}(\tilde{\mathbf{x}})$ is the global field estimated from data samples according to the generalized superposition principle (9).

Instantiation I (Electrostatic Field). An early example of an interaction field is provided by the EFM. Consider a unit point charge of sign ± 1 located at $\tilde{\mathbf{x}}' \in \mathbb{R}^{D+1}$. By Coulomb's law, it induces the field $\mathbf{E}_\pm^{\tilde{\mathbf{x}}'}(\tilde{\mathbf{x}}) \propto \pm(\tilde{\mathbf{x}} - \tilde{\mathbf{x}}')/\|\tilde{\mathbf{x}} - \tilde{\mathbf{x}}'\|_2^{D+1}$. Assigning a positive unit charge to augmented samples $\tilde{\mathbf{x}}_0 \sim \pi_0 \times \delta_{z=0}$ and a negative to $\tilde{\mathbf{x}}_L \sim \pi_L \times \delta_{z=L}$ yields, by the electrostatic superposition principle (Xu et al., 2022; Kolesov et al., 2025), the pairwise-defined field $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}}) = \mathbf{E}_+^{\tilde{\mathbf{x}}_0}(\tilde{\mathbf{x}}) + \mathbf{E}_-^{\tilde{\mathbf{x}}_L}(\tilde{\mathbf{x}})$. This field is an interaction field in the sense of §2.2. Moreover, it is defined on the ambient space $\mathbb{R}^{D+1}$ away from the particle singularities $\tilde{\mathbf{x}}_0$ and $\tilde{\mathbf{x}}_L$, and its field lines may move backward along the z-axis (Manukhov et al., 2026, §2.3). Therefore, the EFM field is not forward-only.

Instantiation II (IFM Field Realization). As discussed above, fields with backward-oriented lines are not straightforward to use in practice. To address this, the authors of IFM propose a forward-only realization. They decompose the field as $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}}) = \|\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})\|_2\, \mathbf{n}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})$, where $\tilde{\mathbf{x}} = (\mathbf{x}, z)$. In (Manukhov et al., 2026, Appendix A.4), they give a concrete procedure to specify $\|\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})\|_2$ and $\mathbf{n}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})$, which we reproduce in Appendix A, for convenience.

2.3 CFM vs. IFM

Previously, we reviewed the CFM and IFM frameworks separately. Although their differences may already be apparent, we include **Table 3** for convenience; it lists the main terms of both frameworks side by side. The key differences can be summarized as follows. First, the IFM **global field** can be estimated by Monte Carlo averaging over particle pairs via the **superposition principle** (9), whereas the CFM **global drift**, defined through a **conditional expectation** (4), is generally intractable to evaluate directly. Second, the training procedures differ: CFM uses a **single-sample conditional velocity as the target** (5) and learns it at points drawn from the **conditional distribution**, while IFM uses the **normalized global field as the target** (14) and learns it at points drawn from heuristic **volume coverage distribution**. Despite these differences, the following example (§2.4) suggests that the two viewpoints may describe the

same underlying dynamics.

2.4 Poisson Flow Generative Models

PFGM is a noise-to-data generative framework rooted in electrostatics (akin to EFM). It is of particular interest here because it admits both a field-based and a flow-based interpretation. To our knowledge, PFGM is the only known framework that supports both viewpoints.

Field-based interpretation. PFGM embeds the data distribution into an extended space $\mathbb{R}^{D+1}$ by placing each sample $\mathbf{x}_0 \sim \pi_0$ on the plane $\{z = 0\}$, i.e., at $\tilde{\mathbf{x}}_0 = (\mathbf{x}_0, 0)$. Each $\tilde{\mathbf{x}}_0$ is treated as a unit positive charge, inducing an electric field in accordance with Coulomb's law:

$$\mathbf{E}^{\mathbf{x}_0}(\tilde{\mathbf{x}}) \propto \frac{\tilde{\mathbf{x}} - \tilde{\mathbf{x}}_0}{\|\tilde{\mathbf{x}} - \tilde{\mathbf{x}}_0\|_2^{D+1}}, \tag{15}$$

for $\tilde{\mathbf{x}} = (\mathbf{x}, z) \in \mathbb{R}^{D+1}$. The global field $\mathbf{E}$ is obtained via the *electrostatic superposition principle*. The field lines of $\mathbf{E}$ then define a transport from the data distribution to a prior π_L, which is a uniform distribution over a hemisphere of radius L (with $L \to \infty$ or sufficiently large in practice).

Remark. Despite its resemblance to IFM, PFGM is not a concrete IFM instance: its conditional field is one-sided (constructed from a single target point $\mathbf{x}_0$), whereas IFM requires two-sided conditioning.

Flow-based interpretation. Subsequent work (Xu et al., 2023, PFGM++) reveals that PFGM possesses an underlying flow-based structure. Specifically, it demonstrates that PFGM's data-space generative dynamics, as in (12), can be obtained by minimizing the objective

$$\mathbb{E}_{z,\mathbf{x}_0,\mathbf{x}_z} \left\| f_\theta(\mathbf{x}_z, z) - \frac{\mathbf{x}_z - \mathbf{x}_0}{z} \right\|_2^2, \tag{16}$$

where $z \sim \mathcal{U}[0, L]$ and $\mathbf{x}_z \sim p_z^{\mathbf{x}_0}(\mathbf{x}) \propto (\|\mathbf{x} - \mathbf{x}_0\|_2^2 + z^2)^{-\frac{D+1}{2}}$. The similarity between (16) and the CFM objective (5) suggests a flow-based interpretation of PFGM: $p_z^{\mathbf{x}_0}$ in (16) acts as a conditional path distribution, and $(\mathbf{x}_z - \mathbf{x}_0)/z$ serves as the corresponding conditional velocity target.

Summary. The PFGM example shows that the same generative transport can be described either as a *flow* or as *field lines*. This motivates two questions:

(1) Is there a general CFM–IFM duality?

(2) Can IFM admit a one-sided (target-only) formulation, analogous to one-sided CFM?

3 Duality of CFM and forward-only IFM

In this section, we formally establish the duality between CFM and forward-only IFM. We first introduce in §3.1 a one-sided formulation of IFM, paralleling the one-sided setting in CFM and enabling a unified treatment of both one-sided and two-sided constructions. We then prove the two directions of the equivalence. In §3.2, we show how any CFM construction induces a forward-only interaction field with the same data-space dynamics. Conversely, in §3.3, we show how any forward-only interaction field can be normalized into a conditional probability path and velocity field, thereby recovering a CFM representation. Finally, in §4, we summarize the theoretical consequences of this equivalence. All proofs are provided in Appendix B.

Before stating our main results, we make a simple observation. In forward-only fields, the spatial coordinate z (the distance from the left capacitor plate) plays the same role as the time variable t in conditional flows. Likewise, the plate separation L corresponds to the terminal time T. With this correspondence in mind, we will use the (t, T)-based notation for both conditional flows and forward-only interaction fields throughout the remainder of the paper.

3.1 One-sided IFM

Standard IFM is defined in a two-sided form. To state the duality for both one-sided and two-sided constructions, we begin by introducing a one-sided notion of an interaction field. Following PFGM (§2.4) and the general IFM formulation (§2.2), we define the one-sided interaction field on the extended space $\mathbb{R}^{D+1}$, where the data distribution lies on a hyperplane $\{t = 0\}$.

Definition (One-sided Interaction Field). A vector field $\mathbf{E}^{\mathbf{x}_0}(\tilde{\mathbf{x}}) \in \mathbb{R}^{D+1}$ is called a one-sided interaction field if it satisfies the following three properties. First, its field lines start at the particle $\tilde{\mathbf{x}}_0 = (\mathbf{x}_0, 0)$: $d\tilde{\mathbf{x}}(\tau)/d\tau = \mathbf{E}^{\mathbf{x}_0}(\tilde{\mathbf{x}}(\tau)), \tilde{\mathbf{x}}(\tau_s) = \tilde{\mathbf{x}}_0$, where τ_s denotes the starting value of the field-line parameter and there exists $\tau_f > \tau_s$ such that $\tilde{x}(\tau_f) \in \mathbb{R}^D \times \{T\}$. Second, the field is divergence-free: $\nabla \cdot \mathbf{E}^{\mathbf{x}_0}(\tilde{\mathbf{x}}) = 0$ for all $\tilde{\mathbf{x}} \in \mathbb{R}^{D+1} \setminus \{\tilde{\mathbf{x}}_0\}$. Third, the field has a fixed total flux.

Remark. This definition matches the two-sided Interaction Field, except that the field lines are not required to end at any predetermined particle.

Generative dynamics. The global field induced by the distribution π_0 is obtained via generalized superposition, analogously to (9): $\mathbf{E}(\tilde{\mathbf{x}}) = \int \mathbf{E}^{\mathbf{x}_0}(\tilde{\mathbf{x}})\, \pi_0(\mathbf{x}_0)\, d\mathbf{x}_0$. The corresponding field lines (10) push forward π_0 to a distribution π_T. Consequently, one-sided IFM defines a source (prior) distribution π_T and transports between π_T and π_0 along the global field lines.

Remark. In contrast to two-sided IFM, the generative dynamics in one-sided IFM is essentially built into the construction, since the source (prior) distribution π_T is defined as the pushforward induced by the field lines. In the two-

sided setting, both endpoint distributions are specified in advance, and the fact that the field lines induce transport between them is a nontrivial property.

Example (PFGM $\in$ forward-only one-sided IFM). The PFGM field defined in §2.4 is an instance of one-sided IFM. Moreover, it is forward-only, i.e, $\mathbf{E}^{\mathbf{x}_0}(\tilde{\mathbf{x}})_t > 0$ for all $t > 0$.

Notation. As in CFM, we use blue to indicate that the source-conditioning point is omitted, corresponding to the one-sided case. Throughout the remainder of this section, $\mathbf{x}_0, \mathbf{x}_T$ denotes either $\mathbf{x}_0$ in the one-sided case or $(\mathbf{x}_0, \mathbf{x}_T)$ in the two-sided case.

3.2 From Flow to Fields

We first establish the flow-to-field direction of the duality.

Theorem 3.1. (CFM $\subseteq$ forward-only IFM). *Let π_0 and π_T be D-dimensional target and source distributions, respectively. Let $(\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}), p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}))$ be a Conditional Flow (§2.1), and let $\mathbf{v}_t(\mathbf{x})$ and $p_t(\mathbf{x})$ denote the global velocity (4) and marginal probability path (3) obtained by marginalizing over $\pi_{0,T}(\mathbf{x}_0, \mathbf{x}_T)$, respectively. Define the $(D+1)$-dimensional vector field*

$$\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t) := (\underbrace{\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})}_{\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_{\mathbf{x}}}, \underbrace{p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})}_{\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t}). \quad (17)$$

Then, under mild assumptions specified in Appendix B.1:

1. *The vector field $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}$ is a forward-only Interaction Field (§2.2) with unit total flux $\Phi_0 = 1$.*
2. *The global field $\mathbf{E}$ obtained by the generalized superposition principle (9) can be recovered as*
$$\mathbf{E}(\mathbf{x},t) = \big(\mathbf{v}_t(\mathbf{x})\, p_t(\mathbf{x}),\ p_t(\mathbf{x})\big). \quad (18)$$
3. *The CFM generative dynamics (4) coincide with the dynamics induced by field lines (10), reparameterized by the distance from the plate (12):*
$$\frac{d\mathbf{x}}{dt} = \mathbf{v}_t(\mathbf{x}) = \frac{\mathbf{E}(\mathbf{x},t)_{\mathbf{x}}}{\mathbf{E}(\mathbf{x},t)_t}. \quad (19)$$

Theorem 3.1 shows that conditional flow admits a forward-only field representation with the same data-space generative dynamics. In this representation, the CFM probability current becomes the spatial component of the interaction field, and the marginal density becomes its time component. Concrete field realizations of common flow-based approaches, including stochastic interpolants, diffusion models, and flow matching, are summarized in Appendix A, Table 4.

3.3 From Fields to Flows

We now establish the converse field-to-flow direction of the duality. Motivated by Theorem 3.1, a natural attempt to recover a conditional flow from a given forward-only interaction field $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)$ would be to set

$$\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \frac{\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_{\mathbf{x}}}{\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t},\ p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t. \quad (20)$$

However, the last component $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t$ of an interaction field is not necessarily normalized:

$$\Phi_t^{\mathbf{x}_0,\mathbf{x}_T} := \int \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t\, d\mathbf{x} \neq 1. \quad (21)$$

Moreover, $\Phi_t^{\mathbf{x}_0,\mathbf{x}_T}$ is a scalar that may, a priori, depend on $\mathbf{x}_0, \mathbf{x}_T$ and t. Nevertheless, it is in fact independent of both, as formalized in our following proposition.

Proposition 3.2. (Flux Conservation across Slices). *Under mild assumptions specified in Appendix B.3 the scalar $\Phi_t^{\mathbf{x}_0,\mathbf{x}_T}$ in (21) is the flux through the slice $\{t = \text{const}\}$ at distance t from the left plate and coincides with the total flux Φ_0 (8) of the forward-only Interaction Field.*

Having established this proposition, we now formalize the field-to-flow direction of the duality.

Theorem 3.3. (forward-only IFM $\subseteq$ CFM). *Let π_0 and π_T be D-dimensional target and source distributions, respectively. Let $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)$ be a forward-only Interaction Field, and let $\mathbf{E}(\mathbf{x},t)$ denote the global field obtained by the generalized superposition principle (9). Define the conditional velocity and conditional density path by*

$$\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \frac{\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_{\mathbf{x}}}{\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t},\ p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \frac{\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t}{\Phi_0}. \quad (22)$$

Then, under mild assumptions specified in Appendix B.3:

1. *The pair $(\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}, p_t^{\mathbf{x}_0,\mathbf{x}_T})$ is a Conditional Flow (§2.1).*
2. *The global drift $\mathbf{v}_t$ (4) and probability path p_t (3) can be recovered from the global field $\mathbf{E}$ via*
$$\mathbf{v}_t(\mathbf{x}) = \frac{\mathbf{E}(\mathbf{x},t)_{\mathbf{x}}}{\mathbf{E}(\mathbf{x},t)_t}, \quad p_t(\mathbf{x}) = \frac{\mathbf{E}(\mathbf{x},t)_t}{\Phi_0}. \quad (23)$$
3. *The CFM generative dynamics (4) coincide with the dynamics induced by field lines (10), reparameterized by the distance from the plate (12):*
$$\frac{d\mathbf{x}}{dt} = \mathbf{v}_t(\mathbf{x}) = \frac{\mathbf{E}(\mathbf{x},t)_{\mathbf{x}}}{\mathbf{E}(\mathbf{x},t)_t}. \quad (24)$$

Our Theorem 3.3 shows that forward-only interaction field admits a conditional flow representation with the same data-space dynamics. The key step is the normalization of the field's time component by the conserved flux Φ_0, which turns the slice flux density into a conditional probability path.

Example (PFGM's flow-based construction). Previously, we showed that PFGM is a forward-only one-sided IFM instance. Therefore, by Theorem 3.3, it admits an equivalent conditional flow representation. This construction recovers the objects introduced in §2.4 as the flow-based interpretation of PFGM. In this sense, the flow-based formulation

used in PFGM++ arises naturally as a special case of the field-to-flow construction.

Example (IFM's flow-based construction). In §2.2, the canonical IFM field is 'constructed directly as a forward-only interaction field. Using our Theorem 3.3, we can express it as a conditional flow, which has a simpler form: $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}_t) = \dot{\mathcal{I}}(t,\mathbf{x}_0,\mathbf{x}_T) + (\dot{\sigma}(t)/\sigma(t))\big(\mathbf{x}_t - \mathcal{I}(t,\mathbf{x}_0,\mathbf{x}_T)\big)$, $p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \mathcal{N}(\mathbf{x} \mid \mathcal{I}(t,\mathbf{x}_0,\mathbf{x}_T), \sigma(t)^2 I)$, where $\mathcal{I}(t,\mathbf{x}_0,\mathbf{x}_T)$ and $\sigma(t)$ are specified in Appendix A, Table 4. This representation makes explicit that the IFM field is induced by the two-sided stochastic interpolant from §2.1, namely, $\mathbf{x}_t = \mathcal{I}(t,\mathbf{x}_0,\mathbf{x}_T) + \sigma(t)\epsilon$. Additional field–flow construction examples are summarized in Appendix A, Table 4.

4 Duality Takeaways and Empirical Verification

We now summarize the main implications of the flow–field duality and support them empirically in a controlled toy setting. All empirical checks use the same Gaussian-to-Gaussian-mixture dataset and two two-sided stochastic interpolants $\{\mathbf{x}_{i,t}\}_{i=1}^2$, defined in Appendix C. The section is organized around three takeaways: for each one, we state the implication of the duality, explain its meaning, and provide supporting empirical evidence. Additional discussions are provided in Appendix D.

Takeaway 1 (Expressiveness). At the level of induced data-space generative dynamics,

$$\text{CFM} \equiv \text{forward-only IFM} \subsetneq \text{IFM}.$$

Theorems 3.1 and 3.3 show that CFM and forward-only IFM induce exactly the same class of data-space generative dynamics. Thus, the probabilistic CFM view and the geometric IFM view are not separate model classes in this regime: they are two representations of the same dynamics.

The strict inclusion comes from the field perspective. General IFM also allows interaction fields with backward-oriented field lines, such as EFM in Figure 1(d). These fields cannot be represented by standard CFM dynamics, and therefore provide a concrete mechanism by which IFM goes beyond conventional conditional flow matching.

This equivalence is constructive. Equations (17) and (22) give explicit maps between conditional flows and forward-only interaction fields, which makes the duality directly testable. For each interpolant $i \in \{1,2\}$, we construct the corresponding conditional interaction field $\mathbf{E}_i^{\mathbf{x}_0,\mathbf{x}_T}$ and compare two independently learned data-space drifts: the CFM drift $\mathbf{v}_i$, trained with the CFM objective (5), and the IFM-induced drift $\mathbf{u}_i(\mathbf{x},t) := \frac{\mathbf{E}_i(\mathbf{x},t)_\mathbf{x}}{\mathbf{E}_i(\mathbf{x},t)_t}$, trained through the normalized-field IFM objective (14).

Table 1 verifies this prediction. Matched CFM–IFM pairs have much smaller error than mismatched pairs: $\mathbf{v}_1$ agrees with $\mathbf{u}_1$, and $\mathbf{v}_2$ agrees with $\mathbf{u}_2$, whereas dynamics built from different interpolants remain far apart. This supports the constructive duality: the two representations coincide when they come from the same interpolant, but different interpolants define genuinely different data-space dynamics. The small residual error in the matched comparisons is attributable to finite-sample estimation, numerical approximation, and imperfect neural-network training.

Type	Comparison	MSE (↓)
Matched	$\mathbf{v}_1$ vs. $\mathbf{u}_1$	**0.052±0.005**
Matched	$\mathbf{v}_2$ vs. $\mathbf{u}_2$	**0.083±0.003**
Mismatched	$\mathbf{v}_1$ vs. $\mathbf{v}_2$	0.479±0.003
Mismatched	$\mathbf{u}_1$ vs. $\mathbf{u}_2$	0.365±0.001

Table 1. Empirical verification of the CFM–forward-only IFM duality. Matched CFM–IFM pairs $(\mathbf{v}_i, \mathbf{u}_i)$, constructed from the same interpolant, have substantially lower MSE than mismatched pairs. Here $\mathbf{u}_i = \mathbf{E}_{i,\mathbf{x}}/\mathbf{E}_{i,t}$ denotes the IFM-induced drift.

Takeaway 2 (IFM probabilistic interpretation). A global forward-only interaction field $\mathbf{E}(\mathbf{x},t)$ induces not only field-line dynamics, but also an explicit data-space distribution path:

$$p_t(\mathbf{x}) = \frac{\mathbf{E}(\mathbf{x},t)_t}{\Phi_0}.$$

Theorem 3.3 shows that the geometric IFM representation has an equivalent probabilistic view. In a forward-only field, the t-component of the global field is a flux density across the slice at time t; after normalization by the conserved flux Φ_0, it becomes the marginal density $p_t(\mathbf{x})$ of the induced data-space dynamics.

We illustrate this interpretation using the field $\mathbf{E}_1$ constructed from the first interpolant. We evolve particles under the IFM-induced dynamics (11) and record their empirical distributions at several times. We then compare these distributions with heatmaps of the normalized field density $\mathbf{E}_1(\mathbf{x},t)_t/\Phi_0$ at the same times. As shown in Figure 2, the transported particles concentrate in the regions of high field density, supporting the probabilistic interpretation predicted by the flow–field duality.

Takeaway 3 (Field-informed volume coverage). Forward-only IFM induces a natural sampling distribution over the augmented space: sample $t \sim \mathcal{U}[0,T]$ and then

$$\mathbf{x} \sim p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}), \qquad p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \frac{\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t}{\Phi_0}.$$

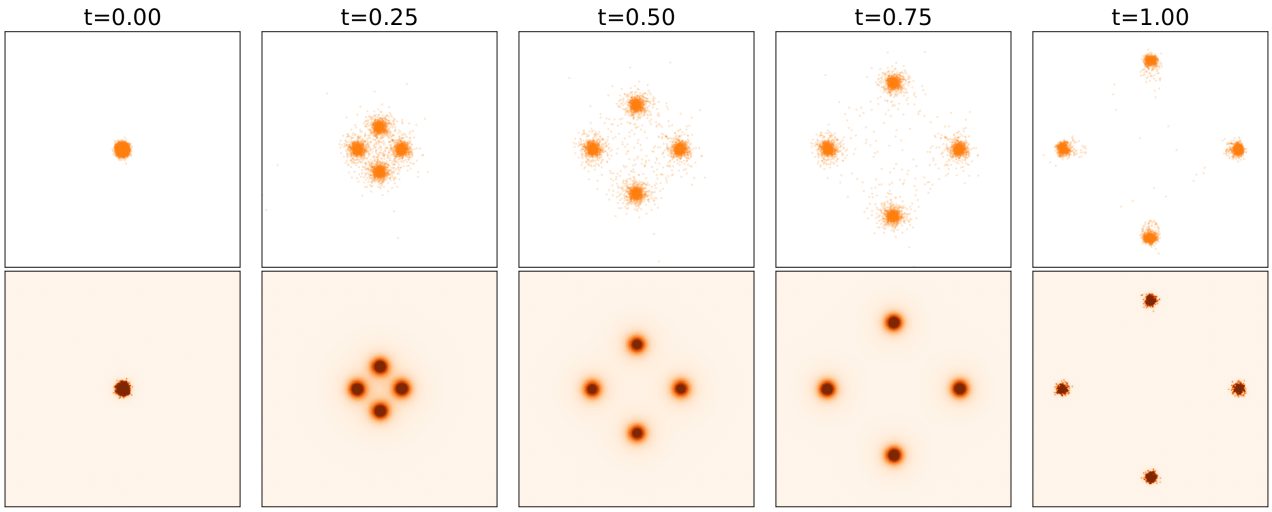

Figure 2. Probabilistic interpretation of a forward-only interaction field. Top: particles evolved under the IFM-induced dynamics. Bottom: the corresponding normalized field density $\mathbf{E}(\mathbf{x}, t)_t/\Phi_0$.

The bidirectional constructions in Equations (17) and (22) show that a forward-only conditional field encodes not only a drift, but also a conditional probability path. In particular, the t-component of $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}$, normalized by the conserved flux Φ_0, defines the conditional distribution $p_t^{\mathbf{x}_0,\mathbf{x}_T}$. This gives a field-informed alternative to the heuristic volume distribution p_{vol} used in IFM training.

One consequence is a CFM-like single-sample regression objective for the data-space drift:

$$\mathbb{E}_{t,\mathbf{x}_0,\mathbf{x}_T,\,\mathbf{x}\sim p_t^{\mathbf{x}_0,\mathbf{x}_T}} \left\| \mathbf{v}_\theta(\mathbf{x},t) - \frac{\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_{\mathbf{x}}}{\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t} \right\|_2^2,$$

where $p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t/\Phi_0$. The same density can also replace the external volume-coverage distribution in the normalized-field IFM objective:

$$\mathbb{E}_{t,\mathbf{x}_0,\mathbf{x}_T,\,\mathbf{x}\sim p_t^{\mathbf{x}_0,\mathbf{x}_T}} \left\| f_\theta(\mathbf{x},t) - \frac{\mathbf{E}(\mathbf{x},t)}{\|\mathbf{E}(\mathbf{x},t)\|_2} \right\|_2^2. \tag{25}$$

To assess whether the field-induced density provides a useful training distribution, we perform a matched–mismatched ablation. For each field $\mathbf{E}_i$, we train objective (25) using either the matched distribution $p_{i,t}^{\mathbf{x}_0,\mathbf{x}_T}$ induced by the same field, or the mismatched distribution $p_{j,t}^{\mathbf{x}_0,\mathbf{x}_T}$ induced by the other field, $j \neq i$. Thus, the objective and target field are fixed, and only the volume-coverage distribution is changed.

Table 2 shows that matched sampling consistently achieves a lower final W_2 distance. This supports the interpretation that the t-component of a forward-only interaction field identifies where training points should be placed in the augmented space, providing a principled field-informed replacement for heuristic volume coverage.

Field	Volume distribution	W_2 (↓)
$\mathbf{E}_1$	**matched** $p_{1,t}^{\mathbf{x}_0,\mathbf{x}_T}$	**0.25±0.04**
	mismatched $p_{2,t}^{\mathbf{x}_0,\mathbf{x}_T}$	0.58±0.03
$\mathbf{E}_2$	mismatched $p_{1,t}^{\mathbf{x}_0,\mathbf{x}_T}$	0.81±0.16
	matched $p_{2,t}^{\mathbf{x}_0,\mathbf{x}_T}$	**0.38±0.04**

Table 2. Ablation of field-induced volume coverage. For each field $\mathbf{E}_i$, sampling from the matched field-induced distribution $p_{i,t}^{\mathbf{x}_0,\mathbf{x}_T}$ yields a lower final W_2 distance than sampling from the mismatched distribution $p_{j,t}^{\mathbf{x}_0,\mathbf{x}_T}$, $j \neq i$.

Impact Statements

This paper presents work whose goal is to advance the field of machine learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

References

Albergo, M. S. and Vanden-Eijnden, E. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022.

Albergo, M. S., Boffi, N. M., and Vanden-Eijnden, E. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023a.

Albergo, M. S., Goldstein, M., Boffi, N. M., Ranganath, R., and Vanden-Eijnden, E. Stochastic interpolants with data-dependent couplings. *arXiv preprint arXiv:2310.03725*, 2023b.

Bogachev, V. I., Röckner, M., and Shaposhnikov, S. V. On

the ambrosio–figalli–trevisan superposition principle for probability solutions to fokker–planck–kolmogorov equations. *Journal of Dynamics and Differential Equations*, 33(2):715–739, 2021.

Gao, R., Hoogeboom, E., Heek, J., De Bortoli, V., Murphy, K. P., and Salimans, T. Diffusion models and gaussian flow matching: Two sides of the same coin. In *The Fourth Blogpost Track at ICLR 2025*, 2025.

Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium, 2018. URL https://arxiv.org/abs/1706.08500.

Karras, T., Laine, S., and Aila, T. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4401–4410, June 2019.

Kolesov, A., Manukhov, S., Palyulin, V. V., and Korotin, A. Field matching: an electrostatic paradigm to generate and transfer data. 2025. URL https://openreview.net/forum?id=9dHilxylvC.

Krizhevsky, A. Learning multiple layers of features from tiny images. Technical report, University of Toronto, April 2009. URL https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf.

Lipman, Y., Chen, R. T., Ben-Hamu, H., Nickel, M., and Le, M. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

Liu, X., Gong, C., and Liu, Q. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.

Liu, Z., Luo, D., Xu, Y., Jaakkola, T., and Tegmark, M. Genphys: From physical processes to generative models. *arXiv preprint arXiv:2304.02637*, 2023.

Manukhov, S. I., Kolesov, A., Palyulin, V. V., and Korotin, A. Interaction field matching: Overcoming limitations of electrostatic models. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=GEsTLuJy1q.

Shlenskii, D. and Korotin, A. Overclocking electrostatic generative models. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=flo449mncA.

Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., and Ganguli, S. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pp. 2256–2265. pmlr, 2015.

Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.

Tong, A., Fatras, K., Malkin, N., Huguet, G., Zhang, Y., Rector-Brooks, J., Wolf, G., and Bengio, Y. Improving and generalizing flow-based generative models with minibatch optimal transport. *arXiv preprint arXiv:2302.00482*, 2023.

Xu, Y., Liu, Z., Tegmark, M., and Jaakkola, T. Poisson flow generative models. *Advances in Neural Information Processing Systems*, 35:16782–16795, 2022.

Xu, Y., Liu, Z., Tian, Y., Tong, S., Tegmark, M., and Jaakkola, T. Pfgm++: Unlocking the potential of physics-inspired generative models. In *International Conference on Machine Learning*, pp. 38566–38591. PMLR, 2023.

Table 3. **CFM vs. IFM**: side-by-side comparison of the Conditional Flow Matching and Interaction Field Matching frameworks.

Aspect	Conditional Flow Matching	Interaction Field Matching
Elementary object	**Conditional Flow** (§2.1) $(\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}),\ p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}))$ defined **in the data space** $\mathbb{R}^D$	**Interaction Field** (6) $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})$ defined **in the extended space** $\mathbb{R}^{D+1}$
Global object	**Conditional averaging** of conditional velocities: $\mathbf{v}_t(\mathbf{x}) = \int \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) p_t(\mathbf{x}_0,\mathbf{x}_T \mid \mathbf{x}_t=\mathbf{x}) d\mathbf{x}_0 d\mathbf{x}_T$ $\neq \int \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, \pi_{0,T}(\mathbf{x}_0,\mathbf{x}_T)\, d\mathbf{x}_0\, d\mathbf{x}_T$	**Generalized superposition** of pairwise fields: $\mathbf{E}(\tilde{\mathbf{x}}) = \int \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})\, \pi_{0,L}(\mathbf{x}_0,\mathbf{x}_L)\, d\mathbf{x}_0\, d\mathbf{x}_L$
Generative dynamics	ODE-based dynamics **in the data space**: $\frac{d\mathbf{x}}{dt} = \mathbf{v}_t(\mathbf{x}), \quad \mathbf{x}_0 \sim \pi_0$	Field-line transport **in the extended space**: $\frac{d\tilde{\mathbf{x}}}{d\tau} = \mathbf{E}(\tilde{\mathbf{x}}), \quad \tilde{\mathbf{x}} \sim \pi_0 \times \delta_{z=0}$ **Forward-only** fields dynamics **in the data space**: $\frac{d\mathbf{x}}{dz} = \frac{\mathbf{E}(\tilde{\mathbf{x}})_{\mathbf{x}}}{\mathbf{E}(\tilde{\mathbf{x}})_z}, \quad \mathbf{x}_0 \sim \pi_0$
Marginal distributions	**Explicitly** defined via conditional distributions: $p_t(\mathbf{x}) = \int p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, \pi_{0,T}(\mathbf{x}_0,\mathbf{x}_T)\, d\mathbf{x}_0\, d\mathbf{x}_T$	**No explicit** marginal density path: defined implicitly via field-line transport (10)
Training objective	**Single-sample conditional velocity** regression: $\mathbb{E}_{t,\mathbf{x}_0,\mathbf{x}_T,\mathbf{x}_t} \left\| \mathbf{v}_t(\mathbf{x}_t) - \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}_t) \right\|_2^2$	**Multi-sample global field** regression: $\mathbb{E}_{\tilde{\mathbf{x}}\sim p_{\text{vol}}} \|f_\theta(\tilde{\mathbf{x}}) - \mathbf{E}(\tilde{\mathbf{x}})\|_2^2$ **Normalized field** in the forward-only case: $\mathbb{E}_{\tilde{\mathbf{x}}\sim p_{\text{vol}}} \left\| f_\theta(\tilde{\mathbf{x}}) - \frac{\mathbf{E}(\tilde{\mathbf{x}})}{\|\mathbf{E}(\tilde{\mathbf{x}})\|_2} \right\|_2^2$
Training volume	**Explicitly** defined by conditional distributions: $\mathbf{x}_0,\mathbf{x}_T \sim \pi_{0,T},\ t \sim \mathcal{U}[0,T],\ \mathbf{x}_t \sim p_t^{\mathbf{x}_0,\mathbf{x}_T}$	**Ad hoc heuristic** distribution: $\tilde{\mathbf{x}} \sim p_{\text{vol}}$

A CFM and IFM construction examples

A.1 Flow-field Dual Constructions

In this section, we provide additional examples of common frameworks formulated via the flow-field dual construction, as summarized in Table 4.

A.2 Original IFM field realization

For completeness, we restate the forward-only Interaction Field realization proposed in Manukhov et al. (2026, Appendix A.4). The construction specifies the magnitude $\|\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})\|_2$ and the unit direction $\mathbf{n}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})$ of the conditional field $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})$, where $\tilde{\mathbf{x}} = (\mathbf{x}, z)$. Concretely, they define

$$\|\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})\|_2 := \frac{\exp\left(-\frac{r_\perp(\tilde{\mathbf{x}})^2}{2\sigma(z)^2}\right)}{\sigma(z)^D \cos\alpha(\tilde{\mathbf{x}})}, \tag{26}$$

$$\mathbf{n}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}}) := \mathbf{e}_\perp \cdot \sin\alpha(\tilde{\mathbf{x}}) + \mathbf{e}'_z \cdot \cos\alpha(\tilde{\mathbf{x}}), \tag{27}$$

$$\mathbf{e}_\perp := \frac{\mathbf{x}_\perp(\tilde{\mathbf{x}})}{r_\perp(\tilde{\mathbf{x}})}, \qquad \mathbf{e}'_z := \frac{\tilde{\mathbf{x}}_L - \tilde{\mathbf{x}}_0}{\|\tilde{\mathbf{x}}_L - \tilde{\mathbf{x}}_0\|_2}, \tag{28}$$

where $\tilde{\mathbf{x}}_0 = (\mathbf{x}_0, 0)$ and $\tilde{\mathbf{x}}_L = (\mathbf{x}_L, L)$ denote the endpoints, and

$$\sigma(z) := \sin(2\pi z/L)$$

(for $d = L/2$ in their notation). The transverse component is defined as

$$\mathbf{x}_\perp(\tilde{\mathbf{x}}) := \mathbf{x} - \mathbf{x}_0(1 - z/L) - \mathbf{x}_L(z/L), \qquad r_\perp(\tilde{\mathbf{x}}) := \|\mathbf{x}_\perp(\tilde{\mathbf{x}})\|_2.$$

Finally, the angle $\alpha(\tilde{\mathbf{x}})$ is chosen so that the ratio $r_\perp(\tilde{\mathbf{x}})/\sigma(z)$ remains constant along each field line, which yields

$$\alpha(\tilde{\mathbf{x}}) := \arctan\left(\frac{2\pi}{L}\, r_\perp(\mathbf{x}, z)\, \cot\left(\frac{2\pi}{L} z\right)\right).$$

Table 4. **Flow-field dual constructions**. This table shows how a common framework can be formulated using both CFM and IFM.

Framework	CFM structure $(\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}),\ p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}))$	IFM structure $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\tilde{\mathbf{x}}) = (\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\tilde{\mathbf{x}})_{\mathbf{x}}, \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\tilde{\mathbf{x}})_t)$, $\tilde{\mathbf{x}} = (\mathbf{x}, t)$
Flow Matching (Lipman et al., 2022)	$\mathbf{v}_t^{\mathbf{x}_0}(\mathbf{x}) = \dfrac{\mathbf{x}-\mathbf{x}_0}{t}$ $p_t^{\mathbf{x}_0}(\mathbf{x}) = \mathcal{N}(\mathbf{x} \mid (1-t)\mathbf{x}_0,\ t^2 I)$	$\mathbf{E}^{\mathbf{x}_0}(\mathbf{x},t)_{\mathbf{x}} = \mathbf{v}_t^{\mathbf{x}_0}(\mathbf{x})\, p_t^{\mathbf{x}_0}(\mathbf{x})$ $\mathbf{E}^{\mathbf{x}_0}(\mathbf{x},t)_t = p_t^{\mathbf{x}_0}(\mathbf{x})$
VE-diffusion (Song et al., 2020)	$\mathbf{v}_t^{\mathbf{x}_0}(\mathbf{x}) = \mathbf{x}-\mathbf{x}_0$ $p_t^{\mathbf{x}_0}(\mathbf{x}) = \mathcal{N}(\mathbf{x} \mid \mathbf{x}_0,\ t^2 I)$	$\mathbf{E}^{\mathbf{x}_0}(\mathbf{x},t)_{\mathbf{x}} = \mathbf{v}_t^{\mathbf{x}_0}(\mathbf{x})\, p_t^{\mathbf{x}_0}(\mathbf{x})$ $\mathbf{E}^{\mathbf{x}_0}(\mathbf{x},t)_t = p_t^{\mathbf{x}_0}(\mathbf{x})$
Linear Interpolant (Albergo et al., 2023a)	$\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \dfrac{\dot\sigma(t)}{\sigma(t)}\left(\mathbf{x}-\mathcal{I}^{\mathbf{x}_0,\mathbf{x}_T}(t)\right)+\dot{\mathcal{I}}^{\mathbf{x}_0,\mathbf{x}_T}(t)$ $p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \mathcal{N}(x \mid \mathcal{I}(t,\mathbf{x}_0,\mathbf{x}_T), \sigma(t)^2 I)$ $\mathcal{I}^{\mathbf{x}_0,\mathbf{x}_T}(t) = \mathbf{x}_0\cdot(1-t/T)+\mathbf{x}_T\cdot t/T$ $\sigma(t) = \sqrt{2\cdot t/T\cdot(1-t/T)}$	$\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_{\mathbf{x}} = \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$ $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t = p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$
PFGM (Xu et al., 2022)	$\mathbf{v}_t^{\mathbf{x}_0}(\mathbf{x}) = \dfrac{\mathbf{x}-\mathbf{x}_0}{t}$ $p_t^{\mathbf{x}_0}(\mathbf{x}) \propto t/\left(\|\mathbf{x}-\mathbf{x}_0\|+t^2\right)^{\frac{D+1}{2}}$	$\mathbf{E}^{\mathbf{x}_0}(\tilde{\mathbf{x}}) = \dfrac{\tilde{\mathbf{x}}-\tilde{\mathbf{x}}_0}{\|\tilde{\mathbf{x}}-\tilde{\mathbf{x}}_0\|_2^{D+1}}$
PFGM++ ($d \geq 1$) (Xu et al., 2023)	$\mathbf{v}_t^{\mathbf{x}_0}(\mathbf{x}) = \dfrac{\mathbf{x}-\mathbf{x}_0}{t}$ $p_t^{\mathbf{x}_0}(\mathbf{x}) \propto t/\left(\|\mathbf{x}-\mathbf{x}_0\|+t^2\right)^{\frac{D+d}{2}}$	$\mathbf{E}^{\mathbf{x}_0}(\mathbf{x},t)_{\mathbf{x}} = \mathbf{v}_t^{\mathbf{x}_0}(\mathbf{x})\, p_t^{\mathbf{x}_0}(\mathbf{x})$ $\mathbf{E}^{\mathbf{x}_0}(\mathbf{x},t)_t = p_t^{\mathbf{x}_0}(\mathbf{x})$
IFM (Manukhov et al., 2026)	$\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \dfrac{\dot\sigma(t)}{\sigma(t)}\left(\mathbf{x}-\mathcal{I}^{\mathbf{x}_0,\mathbf{x}_T}(t)\right)+\dot{\mathcal{I}}^{\mathbf{x}_0,\mathbf{x}_T}(t)$ $p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \mathcal{N}(x \mid \mathcal{I}(t,\mathbf{x}_0,\mathbf{x}_T), \sigma(t)^2 I)$ $\mathcal{I}^{\mathbf{x}_0,\mathbf{x}_T}(t) = \mathbf{x}_0\cdot(1-t/T)+\mathbf{x}_T\cdot t/T$ $\sigma(t) = \sin(2\pi t/T)$	$\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_{\mathbf{x}} = \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$ $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t = p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$ OR originally defined as in §2.2
EFM (Kolesov et al., 2025)	×	$\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\tilde{\mathbf{x}}) = \dfrac{\tilde{\mathbf{x}}-\tilde{\mathbf{x}}_0}{\|\tilde{\mathbf{x}}-\tilde{\mathbf{x}}_0\|_2^{D+1}} - \dfrac{\tilde{\mathbf{x}}-\tilde{\mathbf{x}}_T}{\|\tilde{\mathbf{x}}-\tilde{\mathbf{x}}_T\|_2^{D+1}}$

B Proofs

In this section we provide proofs for all theorems and propositions from the main part of the paper.

B.1 Assumptions for Theorem 3.1

For all $t \in (0,T)$, $\mathbf{x} \in \mathbb{R}^D$:

1. $p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) \in C^1((0,T)\times\mathbb{R}^D)$ is strictly positive;
2. All integrals appearing in the proof (those involving $p_t^{\mathbf{x}_0,\mathbf{x}_T}$ and $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T} p_t^{\mathbf{x}_0,\mathbf{x}_T}$) converge;
3. $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$ is locally Lipschitz continuous in $\mathbf{x}$, uniformly in t on compact subsets of $(0,T)\times\mathbb{R}^D$.

B.2 Proof of Theorem 3.1

We prove the three items in order.

1. $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)$ is a forward-only Interaction Field with $\Phi_0 = 1$.

Forward-only. By definition of the field,

$$\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t = p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) \geq 0.$$

Moreover, if $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t) \neq \mathbf{0}$, then in particular $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t \neq 0$, hence $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t > 0$ at every point where the field is nonzero. Therefore $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}$ is forward-only field.

Field lines connect the two particles. By the superposition principle (Bogachev et al., 2021) for the continuity equation there exist absolutely continuous curves $\mathbf{x}(t)$ satisfying

$$\frac{d\mathbf{x}}{dt} = \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}), \quad \mathbf{x}(0) = \mathbf{x}_0, \quad \mathbf{x}(T) = \mathbf{x}_T.$$

Under the Lipschitz assumption, this ODE with initial conditions has a *unique* solution. Therefore the support of $\mathbf{x}(t)$ consists of a single curve. So the curve $\tilde{\mathbf{x}}(t) := (\mathbf{x}(t), t)$ connects $\tilde{\mathbf{x}}_0 = (\mathbf{x}_0, 0)$ to $\tilde{\mathbf{x}}_T = (\mathbf{x}_T, T)$ in $\mathbb{R}^{D+1}$ and follows

$$\frac{d\tilde{\mathbf{x}}}{dt} = (\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}(t)), 1).$$

Now note that

$$\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t) = \left(\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}),\ p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\right) = p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\,(\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}), 1).$$

Hence $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\tilde{\mathbf{x}})$ is a positive scalar multiple of $(\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}), 1)$ on $p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) > 0$, so it has the same integral curves up to a reparameterization. Concretely, define a new parameter τ along the above curve by

$$\frac{dt}{d\tau} = p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}(t)) \qquad \left(\text{equivalently } \frac{d\tau}{dt} = \frac{1}{p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}(t))}\right),$$

and set $\tilde{\mathbf{x}}(\tau) := \tilde{\mathbf{x}}(t(\tau))$. Then

$$\begin{aligned}
\frac{d\tilde{\mathbf{x}}}{d\tau} &= \frac{d\tilde{\mathbf{x}}}{dt}\frac{dt}{d\tau} \\
&= (\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}(t)), 1)\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}(t)) \\
&= \left(\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}(t)) p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}(t)),\ p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}(t))\right) \\
&= \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\tilde{\mathbf{x}}(\tau)).
\end{aligned}$$

Therefore $\tilde{\mathbf{x}}(\tau)$ is a field line of $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}$ connecting $\tilde{\mathbf{x}}_0$ to $\tilde{\mathbf{x}}_T$.

Divergence-free. Using $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t)_{\mathbf{x}} = \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$ and $p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t)_t$, we obtain

$$\begin{aligned}
\nabla_{(\mathbf{x},t)} \cdot \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t) &= \partial_t \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t)_t + \nabla_{\mathbf{x}} \cdot \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t)_{\mathbf{x}} \\
&= \partial_t p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) + \nabla_{\mathbf{x}} \cdot \left(\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\right) = 0,
\end{aligned}$$

where the last equality is the continuity equation for the Conditional Flow.

Fixed total flux and $\Phi_0 = 1$. Let ∂M be any closed surface that encloses $\tilde{\mathbf{x}}_0 = (\mathbf{x}_0, 0)$ but not $\tilde{\mathbf{x}}_T = (\mathbf{x}_T, T)$. We show that $\int_{\partial M} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} = 1$, which identifies the total flux as $\Phi_0 = 1$ and shows it is independent of $(\mathbf{x}_0, \mathbf{x}_T)$.

First, let's reduce the flux on ∂M to the flux on a small pillbox around $\tilde{\mathbf{x}}_0$. Since $\tilde{\mathbf{x}}_0$ lies in the interior of M, we can choose $R > 0$ and $\epsilon \in (0, T)$ such that the pillbox

$$P := B_R(\mathbf{x}_0) \times [-\epsilon, \epsilon]$$

is contained in M (and in particular does not contain $\tilde{\mathbf{x}}_T$ because $\epsilon < T$). Consider the region $M \setminus P$. It contains neither endpoint, so $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}$ is divergence-free there. Hence, by the divergence theorem,

$$0 = \int_{M\setminus P} \nabla \cdot \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} = \int_{\partial(M\setminus P)} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} = \int_{\partial M} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} - \int_{\partial P} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S},$$

where the minus sign comes from the fact that the normal on ∂P is inward for the region $M \setminus P$. Therefore,

$$\int_{\partial M} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} = \int_{\partial P} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S}.$$

Next, let's compute the flux through the pillbox ∂P. Decompose ∂P into bottom, lateral, and top parts:

$$\partial P = \underbrace{B_R(\mathbf{x}_0) \times \{-\epsilon\}}_{\text{bottom}} \cup \underbrace{\partial B_R(\mathbf{x}_0) \times [-\epsilon, \epsilon]}_{\text{lateral}} \cup \underbrace{B_R(\mathbf{x}_0) \times \{\epsilon\}}_{\text{top}}.$$

Bottom. By the forward-only convention, $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t) = 0$ for $t < 0$, hence

$$\int_{B_R\times\{-\epsilon\}} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} = -\int_{B_R} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, -\epsilon)_t\, d\mathbf{x} = 0.$$

Lateral. For $t \in (0, \epsilon)$ we have $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t)_{\mathbf{x}} = \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$ and $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t)_t = p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$. Integrate the

continuity equation $\partial_t p_t^{\mathbf{x}_0,\mathbf{x}_T} + \nabla_{\mathbf{x}} \cdot (\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}\, p_t^{\mathbf{x}_0,\mathbf{x}_T}) = 0$ over $B_R(\mathbf{x}_0)$ and apply the divergence theorem in $\mathbf{x}$:

$$\frac{d}{dt}\int_{B_R(\mathbf{x}_0)} p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, d\mathbf{x} = -\int_{\partial B_R(\mathbf{x}_0)} \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) \cdot \mathbf{n}\, dS.$$

Integrating from $t = 0$ to $t = \epsilon$ yields

$$\int_0^\epsilon \int_{\partial B_R(\mathbf{x}_0)} \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) \cdot \mathbf{n}\, dS\, dt = \int_{B_R(\mathbf{x}_0)} p_{t=0}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, d\mathbf{x} - \int_{B_R(\mathbf{x}_0)} p_{t=\epsilon}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, d\mathbf{x}.$$

Since $p_{t=0}^{\mathbf{x}_0,\mathbf{x}_T} = \delta_{\mathbf{x}_0}$ and $\mathbf{x}_0 \in B_R(\mathbf{x}_0)$, the first integral equals 1. Therefore the lateral flux equals

$$\int_{\partial B_R \times [-\epsilon,\epsilon]} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} = \int_0^\epsilon \int_{\partial B_R} \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}\, p_t^{\mathbf{x}_0,\mathbf{x}_T} \cdot \mathbf{n}\, dS\, dt = 1 - \int_{B_R} p_{t=\epsilon}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, d\mathbf{x}.$$

Top. Finally, for the top part

$$\int_{B_R \times \{\epsilon\}} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} = \int_{B_R} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},\epsilon)_t\, d\mathbf{x} = \int_{B_R} p_{t=\epsilon}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, d\mathbf{x}.$$

As a result, adding the fluxes through the bottom, lateral, and top part gives

$$\int_{\partial P} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} = 0 + \left(1 - \int_{B_R} p_{t=\epsilon}^{\mathbf{x}_0,\mathbf{x}_T}\, d\mathbf{x}\right) + \int_{B_R} p_{t=\epsilon}^{\mathbf{x}_0,\mathbf{x}_T}\, d\mathbf{x} = 1.$$

Therefore $\int_{\partial M} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} = 1$ for every closed surface ∂M enclosing $\tilde{\mathbf{x}}_0$ but not $\tilde{\mathbf{x}}_T$. Hence the total flux is fixed and equals $\Phi_0 = 1$, independent of $(\mathbf{x}_0, \mathbf{x}_T)$.

This proves that $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}$ is a forward-only Interaction Field with unit total flux.

2. Recovering the global field $\mathbf{E}(\mathbf{x}, t)$.

By the generalized superposition principle,

$$\mathbf{E}(\mathbf{x}, t) = \int \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t)\, d\pi_{0,T}(\mathbf{x}_0, \mathbf{x}_T).$$

Taking components and using the definition (17) gives

$$\mathbf{E}(\mathbf{x}, t)_t = \int p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, d\pi_{0,T}(\mathbf{x}_0, \mathbf{x}_T) = p_t(\mathbf{x}),$$

where the last equality is exactly the marginalization formula for p_t. Similarly,

$$\mathbf{E}(\mathbf{x}, t)_{\mathbf{x}} = \int \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, d\pi_{0,T}(\mathbf{x}_0, \mathbf{x}_T).$$

By the definition of the CFM global velocity $\mathbf{v}_t(\mathbf{x})$ (conditional expectation given $x_t = \mathbf{x}$),

$$\mathbf{v}_t(\mathbf{x}) = \frac{\int \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, d\pi_{0,T}(\mathbf{x}_0, \mathbf{x}_T)}{\int p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, d\pi_{0,T}(\mathbf{x}_0, \mathbf{x}_T)} = \frac{\mathbf{E}(\mathbf{x}, t)_{\mathbf{x}}}{p_t(\mathbf{x})}.$$

Therefore $\mathbf{E}(\mathbf{x}, t)_{\mathbf{x}} = \mathbf{v}_t(\mathbf{x}) p_t(\mathbf{x})$, and hence

$$\mathbf{E}(\mathbf{x}, t) = \big(\mathbf{v}_t(\mathbf{x})\, p_t(\mathbf{x}),\, p_t(\mathbf{x})\big).$$

3. Field-line dynamics coincide with the CFM dynamics.

For a forward-only global field, the data-space dynamics obtained by reparameterizing field lines by t is

$$\frac{d\mathbf{x}}{dt} = \frac{\mathbf{E}(\mathbf{x}, t)_{\mathbf{x}}}{\mathbf{E}(\mathbf{x}, t)_t},$$

Using Item **2**,

$$\frac{\mathbf{E}(\mathbf{x}, t)_{\mathbf{x}}}{\mathbf{E}(\mathbf{x}, t)_t} = \frac{\mathbf{v}_t(\mathbf{x}) p_t(\mathbf{x})}{p_t(\mathbf{x})} = \mathbf{v}_t(\mathbf{x}),$$

so the induced dynamics coincide with the CFM generative ODE:

$$\frac{d\mathbf{x}}{dt} = \mathbf{v}_t(\mathbf{x}) = \frac{\mathbf{E}(\mathbf{x}, t)_{\mathbf{x}}}{\mathbf{E}(\mathbf{x}, t)_t}.$$

□

Remark B.1. The locally Lipschitz condition on the conditional velocity field $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$ is essential for the proof of Theorem 3.1. It ensures the existence of a unique solution to the ODE $\frac{d\mathbf{x}}{dt} = \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$ with prescribed endpoints, which in turn guarantees that the constructed field lines connect $\tilde{\mathbf{x}}_0$ to $\tilde{\mathbf{x}}_T$ without ambiguity. Without this regularity, the duality between conditional flows and interaction fields may break down due to non-unique or ill-defined trajectories.

B.3 Assumptions for Theorem 3.3

Assume $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)$ is a forward-only Interaction Field with total flux Φ_0, satisfying:

1. $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \in C^1(\mathbb{R}^D \times (0,T))$, $\mathbf{E}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t) > 0$ for $t \in (0,T)$ and $\mathbf{E}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t) = 0$ for $t < 0$;
2. The measures converge $\lim_{t\to 0} \mathbf{E}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t) \to \Phi_0 \delta_{\mathbf{x}_0}(\mathbf{x})$, $\lim_{t\to T} \mathbf{E}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t) = \Phi_0 \delta_{\mathbf{x}_T}(\mathbf{x})$ in the weak sense;
3. $\mathbf{E}_{\mathbf{x}}^{\mathbf{x}_0,\mathbf{x}_T} / \mathbf{E}_t^{\mathbf{x}_0,\mathbf{x}_T}$ is locally Lipschitz in $\mathbf{x}$, uniformly in t on compact subsets of $(0,T) \times \mathbb{R}^D$;
4. The Field decays sufficiently at infinity so that integrals converge ($\int_{\mathbb{R}^D} |\mathbf{E}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)| d\mathbf{x} < \infty$) and the flux through infinite lateral surfaces vanishes:
$$\lim_{R\to\infty} \int_{\partial B_R \times [t_1,t_2]} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t) \cdot d\mathbf{S} = 0.$$

B.4 Proof of Proposition 3.2.

Here we imply the above Assumptions 1 and 4. Fix any $t^* \in (0,T)$. We show that the slice flux equals the total flux:
$$\int_{\mathbb{R}^D} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t^*)_t \, d\mathbf{x} = \Phi_0. \tag{29}$$

First, for $R > 0$, define the cylinder
$$M_R := B_R(\mathbf{x}_0) \times [-t^*, t^*], \qquad B_R(\mathbf{x}_0) := \{\mathbf{x} \in \mathbb{R}^D : \|\mathbf{x} - \mathbf{x}_0\|_2 \le R\}.$$
Since $t^* < T$, the volume M_R encloses $\tilde{\mathbf{x}}_0 = (\mathbf{x}_0, 0)$ but not $\tilde{\mathbf{x}}_T = (\mathbf{x}_T, T)$. Hence, by the defining total-flux property of a forward-only Interaction Field,
$$\int_{\partial M_R} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} = \Phi_0.$$

Next, let's decompose the boundary into bottom, lateral, and top parts:
$$\partial M_R = \underbrace{B_R \times \{-t^*\}}_{\text{bottom}} \cup \underbrace{\partial B_R \times [-t^*, t^*]}_{\text{lateral}} \cup \underbrace{B_R \times \{t^*\}}_{\text{top}}.$$

Using outward normals, we obtain
$$\int_{\partial M_R} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} = -\int_{B_R} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},-t^*)_t \, d\mathbf{x} + \int_{\partial B_R \times [-t^*,t^*]} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} + \int_{B_R} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t^*)_t \, d\mathbf{x}. \tag{30}$$

The bottom term vanishes because (by the forward-only convention) $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t) = \mathbf{0}$ for $t < 0$:
$$\int_{B_R} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},-t^*)_t \, d\mathbf{x} = 0.$$

Use the assumption that the field has vanishing lateral flux at infinity (Manukhov et al., 2026):
$$\lim_{R\to\infty} \int_{\partial B_R \times [-t^*,t^*]} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} = 0.$$

Letting $R \to \infty$ in (30) therefore yields
$$\Phi_0 = \lim_{R\to\infty} \int_{\partial M_R} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} = \lim_{R\to\infty} \int_{B_R} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t^*)_t \, d\mathbf{x} = \int_{\mathbb{R}^D} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t^*)_t \, d\mathbf{x},$$
which proves (29). □

B.5 Proof of Theorem 3.3

We prove the three items in order.

1. **The pair $(\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}), p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}))$ defines a Conditional Flow.**

   Well-definedness. Since $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}$ is forward-only, $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t > 0$ wherever the field is defined. Hence $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$ is well-defined and $p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) \geq 0$.

   Normalization. By Proposition 3.2, for every $t \in (0,T)$,

   $$\int_{\mathbb{R}^D} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t \, d\mathbf{x} = \Phi_0.$$

   Therefore,

   $$\int_{\mathbb{R}^D} p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) \, d\mathbf{x} = \frac{1}{\Phi_0} \int_{\mathbb{R}^D} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t \, d\mathbf{x} = 1.$$

   Continuity equation. Using $p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t/\Phi_0$ and $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_{\mathbf{x}}/\Phi_0$, we obtain

   $$\begin{aligned}\partial_t p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) + \nabla_{\mathbf{x}} \cdot \left(\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\right) &= \frac{1}{\Phi_0}\left(\partial_t \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t + \nabla_{\mathbf{x}} \cdot \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_{\mathbf{x}}\right) \\ &= \frac{1}{\Phi_0} \nabla_{(\mathbf{x},t)} \cdot \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t) = 0,\end{aligned}$$

   where the last equality is the divergence-free property of Interaction Fields.

   Boundary conditions. In the IFM capacitor setup, the conditional field has only the endpoint source/sink on the plates. Equivalently, the plate flux is concentrated at the endpoints (suggested to be stated as part of the forward-only IFM definition / standing assumptions): in the sense of measures (distributions) on $\mathbb{R}^D$,

   $$\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},0)_t = \Phi_0\, \delta_{\mathbf{x}_0}(\mathbf{x}), \qquad \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},T)_t = \Phi_0\, \delta_{\mathbf{x}_T}(\mathbf{x}).$$

   Dividing by Φ_0 gives

   $$p_{t=0}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \delta_{\mathbf{x}_0}(\mathbf{x}), \qquad p_{t=T}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \delta_{\mathbf{x}_T}(\mathbf{x}),$$

   which is exactly (2). Thus $(\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}), p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}))$ is a Conditional Flow.

2. **Recovering $(\mathbf{v}_t(\mathbf{x}), p_t(\mathbf{x}))$ from the global field $\mathbf{E}(\mathbf{x},t)$.**

   By generalized superposition (9),

   $$\mathbf{E}(\mathbf{x},t) \;=\; \int \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)\, \pi_{0,T}(\mathbf{x}_0,\mathbf{x}_T)\, d\mathbf{x}_0\, d\mathbf{x}_T.$$

   Taking the t-component and using $p_t^{\mathbf{x}_0,\mathbf{x}_T} = \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}{}_t/\Phi_0$ gives

   $$\mathbf{E}(\mathbf{x},t)_t = \int \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t\, \pi_{0,T}\, d\mathbf{x}_0\, d\mathbf{x}_T = \Phi_0 \int p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, \pi_{0,T}\, d\mathbf{x}_0\, d\mathbf{x}_T.$$

   By the CFM marginalization formula (3), $p_t(\mathbf{x}) = \int p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, \pi_{0,T}\, d\mathbf{x}_0\, d\mathbf{x}_T$, hence

   $$p_t(\mathbf{x}) = \frac{\mathbf{E}(\mathbf{x},t)_t}{\Phi_0}.$$

   Similarly, using $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_{\mathbf{x}}/\Phi_0$,

   $$\mathbf{E}(\mathbf{x},t)_{\mathbf{x}} = \int \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_{\mathbf{x}}\, \pi_{0,T}\, d\mathbf{x}_0\, d\mathbf{x}_T = \Phi_0 \int \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, \pi_{0,T}\, d\mathbf{x}_0\, d\mathbf{x}_T.$$

   On the other hand, by the definition of the global drift (4),

   $$\begin{aligned}\mathbf{v}_t(\mathbf{x})\, p_t(\mathbf{x}) &= \mathbf{E}_{\mathbf{x}_0,\mathbf{x}_T}\left[\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}_t) \mid \mathbf{x}_t = \mathbf{x}\right] p_t(\mathbf{x}) \\ &= p_t(\mathbf{x}) \int \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t(\mathbf{x}_0,\mathbf{x}_T \mid \mathbf{x}_t = \mathbf{x})\, d\mathbf{x}_0\, d\mathbf{x}_T \\ &= p_t(\mathbf{x}) \int \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, \frac{p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, \pi_{0,T}(\mathbf{x}_0,\mathbf{x}_T)}{p_t(\mathbf{x})}\, d\mathbf{x}_0\, d\mathbf{x}_T \\ &= \int \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, \pi_{0,T}(\mathbf{x}_0,\mathbf{x}_T)\, d\mathbf{x}_0\, d\mathbf{x}_T.\end{aligned}$$

Combining the last three displays yields

$$\mathbf{v}_t(\mathbf{x})\, p_t(\mathbf{x}) = \frac{\mathbf{E}(\mathbf{x}, t)_{\mathbf{x}}}{\Phi_0}.$$

Substituting $p_t(\mathbf{x}) = \mathbf{E}(\mathbf{x}, t)_t / \Phi_0$ gives

$$\mathbf{v}_t(\mathbf{x}) = \frac{\mathbf{E}(\mathbf{x}, t)_{\mathbf{x}}}{\mathbf{E}(\mathbf{x}, t)_t},$$

proving (23).

3. **Flow-based dynamics coincide with global field lines.**

Let $\tilde{\mathbf{x}}(\tau) = (\mathbf{x}(\tau), t(\tau))$ be a field line of the global field:

$$\frac{d\tilde{\mathbf{x}}(\tau)}{d\tau} = \mathbf{E}(\tilde{\mathbf{x}}(\tau)).$$

Since $\mathbf{E}$ is a nonnegative superposition of forward-only fields, $\mathbf{E}_t(\mathbf{x}, t) \geq 0$ everywhere (and $\mathbf{E}_t > 0$ along any nontrivial field line). Hence $\frac{dt}{d\tau} = \mathbf{E}_t(\tilde{\mathbf{x}}(\tau)) > 0$ and we may reparameterize the curve by t:

$$\frac{d\mathbf{x}}{dt} = \frac{d\mathbf{x}/d\tau}{dt/d\tau} = \frac{\mathbf{E}(\mathbf{x}, t)_{\mathbf{x}}}{\mathbf{E}(\mathbf{x}, t)_t} = \mathbf{v}_t(\mathbf{x}),$$

where the last equality follows from item **2**. This is exactly the CFM generative dynamics (4).

□

C Empirical Validation Setup

We describe the common setup used for the empirical validations in Section 4. All experiments are performed on the two-dimensional toy dataset shown in Figure 3.

We use two two-sided stochastic interpolants. For $i \in \{1, 2\}$ and $t \in [0, 1]$, define

$$\mathbf{x}_t^i = \mathcal{I}^i(t, \mathbf{x}_0, \mathbf{x}_1) + \sigma^i(t)\epsilon, \qquad \epsilon \sim \mathcal{N}(0, I_2).$$

Equivalently,

$$p_{i,t}^{\mathbf{x}_0, \mathbf{x}_T}(\mathbf{x}) = \mathcal{N}\left(\mathbf{x} \mid \mathcal{I}^i(t, \mathbf{x}_0, \mathbf{x}_1), (\sigma^i(t))^2 I_2\right).$$

The first interpolant uses a linear path:

$$\mathcal{I}^1(t, \mathbf{x}_0, \mathbf{x}_1) = (1 - t)\mathbf{x}_0 + t\mathbf{x}_1, \qquad \sigma^1(t) = 0.05\sqrt{t(1 - t)}.$$

The second interpolant uses a trigonometric path:

$$\mathcal{I}^2(t, \mathbf{x}_0, \mathbf{x}_1) = \cos(\pi t/2)\mathbf{x}_0 + \sin(\pi t/2)\mathbf{x}_1, \qquad \sigma^2(t) = 0.05\sqrt{t(1 - t)}.$$

Each interpolant induces a conditional flow $(\mathbf{v}_{i,t}^{\mathbf{x}_0, \mathbf{x}_T}, p_{i,t}^{\mathbf{x}_0, \mathbf{x}_T})$, where

$$\mathbf{v}_{i,t}^{\mathbf{x}_0, \mathbf{x}_T}(\mathbf{x}) = \partial_t \mathcal{I}^i(t, \mathbf{x}_0, \mathbf{x}_1) + \frac{\dot{\sigma}^i(t)}{\sigma^i(t)}\left(\mathbf{x} - \mathcal{I}^i(t, \mathbf{x}_0, \mathbf{x}_1)\right).$$

By the flow-to-field construction, it also induces a forward-only conditional interaction field

$$\mathbf{E}_i^{\mathbf{x}_0, \mathbf{x}_T}(\mathbf{x}, t) = \left(\mathbf{v}_{i,t}^{\mathbf{x}_0, \mathbf{x}_T}(\mathbf{x})\, p_{i,t}^{\mathbf{x}_0, \mathbf{x}_T}(\mathbf{x}),\ p_{i,t}^{\mathbf{x}_0, \mathbf{x}_T}(\mathbf{x})\right),$$

with global field obtained by superposition principles:

$$\mathbf{E}_i(\mathbf{x}, t) = \mathbb{E}_{\mathbf{x}_0, \mathbf{x}_T}\left[\mathbf{E}_i^{\mathbf{x}_0, \mathbf{x}_T}(\mathbf{x}, t)\right].$$

These two interpolants define the matched CFM–IFM pairs used throughout Section 4.

For all experiments in Section 4, unless stated otherwise, we train IFM using the normalized-field objective, with the target field estimated from 512 endpoint pairs. In the final volume-swapping ablation, we instead use a single-pair field estimate. This choice reflects the practical high-dimensional regime, where multi-sample field estimation can become essentially single-sample due to estimator degeneracy, as discussed in Appendix D.

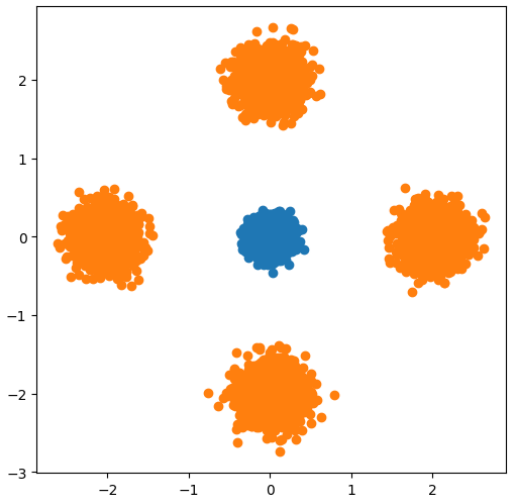

Figure 3. Two-dimensional toy Gaussian dataset used for the empirical validation. Blue points denote samples from the source distribution p_1, a compact Gaussian centered near the origin. Orange points denote samples from the target distribution p_0, a symmetric four-component Gaussian mixture with modes arranged along the coordinate axes.

D Additional Discussion and Experiments on Duality

Rather than introducing a new generative model, we study several widely used methods: the two-sided stochastic interpolant, EDM, and PFGM++, under a unified experimental protocol, focusing on how their behavior changes when trained with multi-sample targets, which is essential for IFM with multi-sample field in target. We build on the official PFGM++ codebase (Xu et al., 2023), preserving the original architecture and training configuration as faithfully as possible, with the only notable modification being a training batch size of $B = 256$ instead of 512 for CIFAR-10 (Krizhevsky, 2009) dataset. Experiments for images were done for unconditional generation task on the CIFAR-10 and FFHQ-64 (Karras et al., 2019) datasets. For CIFAR-10 models were trained for 100k kimg and for 200k kimg on FFHQ, where one kimg corresponds to one thousand training images processed. For multi-sample target computation, we use a target set of size $N \geq B$: when $N > 1$, the target is estimated from the current batch of size B (on which the loss is evaluated) together with $N - B$ additional random samples drawn from the dataset, so the target estimate always uses a pool at least as large as the gradient-evaluation batch. For each FID (Heusel et al., 2018) evaluation, we draw 50k samples and compare against the full training set, using 50 solver steps and the Euler sampler for the two-sided interpolant and the Heun sampler for EDM and PFGM++. All image experiments were done on 2x A100 GPUs, each CIFAR-10 training took around 2 days, and each FFHQ training took around 8 days.

We adopt the velocity formulation and train the network to approximate the multi-sample velocity target $\mathbf{v}_t^N(\mathbf{x})$:

$$\hat{\mathbf{v}}_t^N(\mathbf{x}) := \frac{\hat{\mathbf{E}}^N(\mathbf{x},t)_{\mathbf{x}}}{\hat{\mathbf{E}}^N(\mathbf{x},t)_t} = \frac{\sum_{i=1}^N \mathbf{E}^{\mathbf{x}_0^i,\mathbf{x}_T^i}(\mathbf{x})_{\mathbf{x}}}{\sum_{i=1}^N \mathbf{E}^{\mathbf{x}_0^i,\mathbf{x}_T^i}(\mathbf{x})_t} \tag{31}$$

$$= \frac{\sum_{i=1}^N \mathbf{v}_t^{\mathbf{x}_0^i,\mathbf{x}_T^i}(\mathbf{x})\, p_t^{\mathbf{x}_0^i,\mathbf{x}_T^i}(\mathbf{x})}{\sum_{i=1}^N p_t^{\mathbf{x}_0^i,\mathbf{x}_T^i}(\mathbf{x})} = \sum_{i=1}^N \omega_i\, \mathbf{v}_t^{\mathbf{x}_0^i,\mathbf{x}_T^i}(\mathbf{x}), \tag{32}$$

where $\omega_i := p_t^{\mathbf{x}_0^i,\mathbf{x}_T^i}(\mathbf{x}) / \sum_{k=1}^N p_t^{\mathbf{x}_0^k,\mathbf{x}_T^k}(\mathbf{x})$, $\mathbf{x}_0^i, \mathbf{x}_T^i \sim \pi_{0,T}$, and N is the number of samples used for velocity estimation.

Two-sided interpolant. Following (Albergo & Vanden-Eijnden, 2022), the two-sided interpolant is defined by $\mathbf{x}_t = (1-t)\,\mathbf{x}_0 + t\,\mathbf{x}_T + C\sqrt{t(1-t)}\,\varepsilon$ with $\varepsilon \sim \mathcal{N}(0, I)$, where C controls the injected noise variance. Since the original paper includes an additional $\sqrt{2}$ factor, we report results in terms of $s := C/\sqrt{2}$. Figure 4 shows that larger s consistently degrades sample quality; we use $s = 0.1$ throughout. The model is trained to approximate the conditional velocity $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}$ from Table 4 via Equation (31).

One-sided interpolants. For EDM and PFGM++, the conditional velocity fields are given in Table 4, and a multi-sample velocity estimate $\mathbf{v}_\sigma^N(\mathbf{x}_\sigma)$ can be built via Equation (31) by computing weights from $p_\sigma(\mathbf{x}_\sigma \mid \mathbf{x}_0^i)$ for each $\mathbf{x}_0^i \sim \pi_0$. To preserve the original hyperparameters and sampling procedure, we convert this into a clean-sample estimate:

$$\mathbf{x}_0^N(\mathbf{x}_\sigma) = \mathbf{x}_\sigma - \sigma\, \mathbf{v}_\sigma^N(\mathbf{x}_\sigma) \tag{33}$$

$$= \mathbf{x}_\sigma - \sigma\, \frac{\sum_{i=1}^N \mathbf{v}_\sigma^{\mathbf{x}_0^i}(\mathbf{x}_\sigma)\, p_\sigma(\mathbf{x}_\sigma \mid \mathbf{x}_0^i)}{\sum_{i=1}^N p_\sigma(\mathbf{x}_\sigma \mid \mathbf{x}_0^i)}$$

$$= \mathbf{x}_\sigma - \frac{\sum_{i=1}^N (\mathbf{x}_\sigma - \mathbf{x}_0^i)\, p_\sigma(\mathbf{x}_\sigma \mid \mathbf{x}_0^i)}{\sum_{i=1}^N p_\sigma(\mathbf{x}_\sigma \mid \mathbf{x}_0^i)}, \tag{34}$$

where $\mathbf{x}_0^i \sim \pi_0$ for $i \in \{1, \dots, N\}$. Unlike the two-sided case, these models are trained to predict $\mathbf{x}_0$ directly, following the original EDM and PFGM++ conventions.

Results. Table 5 reports FID as a function of the number of samples N used for target estimation on the CIFAR-10 dataset, while Figure 5 shows training convergence for different target-set sizes N. The observed effect is strongly method-dependent.

For the two-sided interpolant, the training curves are nearly indistinguishable across all tested values of N (Figure 5a), and the FID scores in Table 5 remain almost unchanged, indicating minimal sensitivity to multi-sample targets. This behavior arises because the weight distribution in Eq. (31) is highly concentrated: a single conditional term dominates the sum, and empirically this dominant term is almost always the pair $(\mathbf{x}_0, \mathbf{x}_T)$ used to generate the intermediate point $\mathbf{x}_t \sim p_t^{\mathbf{x}_0, \mathbf{x}_T}$. As a result, adding more samples contributes only negligibly to the estimator. We verified this by removing the dominant pair from the target estimator, which caused the FID score to exceed 100, demonstrating that the remaining samples are insufficient to provide an adequate replacement.

In contrast, for EDM and PFGM++ (Figures 5b–d), increasing N improves both convergence speed and final FID, with consistent quantitative gains shown in Table 5. In these methods, the weight distribution becomes less degenerate as N increases, allowing additional conditional samples to contribute meaningfully to the target estimate. This observation suggests that variance reduction could potentially be improved further through relevance-based sampling of conditional pairs instead of uniform sampling.

Table 5. FID (↓) (Heusel et al., 2018) for CIFAR-10 dataset as a function of the number of samples N used for the multi-sample target. Best scores per column are **bolded**. Standard deviation was measured across 3 different sets of seeds for FID evaluation. Values of standard deviation show that FID improvement is statistically significant for CIFAR-10 dataset.

N	**FID** (↓)			
	Two-sided Linear	**EDM**	**PFGM++** ($d = 128$)	**PFGM++** ($d = 2048$)
1	$\mathbf{3.03 \pm 0.03}$	2.29 ± 0.04	2.40 ± 0.04	2.40 ± 0.05
256	3.04 ± 0.03	2.21 ± 0.03	$\mathbf{2.24 \pm 0.03}$	$\mathbf{2.14 \pm 0.04}$
2048	3.05 ± 0.04	$\mathbf{2.12 \pm 0.02}$	2.28 ± 0.03	2.19 ± 0.03

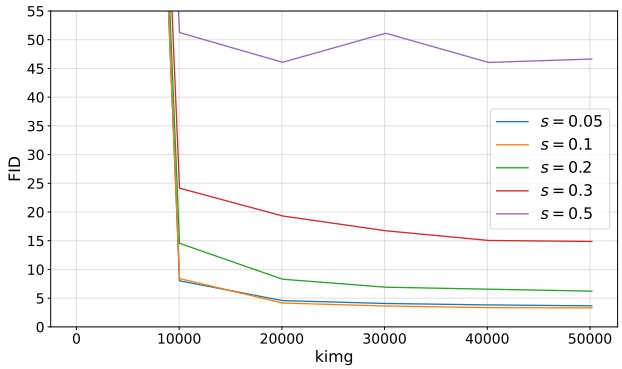

Figure 4. FID vs. training progress for the two-sided interpolant at different noise scales s. Larger s consistently harms sample quality; $s = 0.1$ yields the best results.

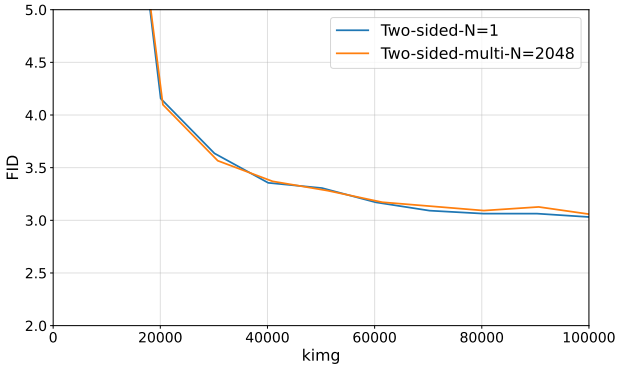

(a) Two-sided interpolant ($s = 0.1$). FID curves are nearly identical across N, showing minimal sensitivity.

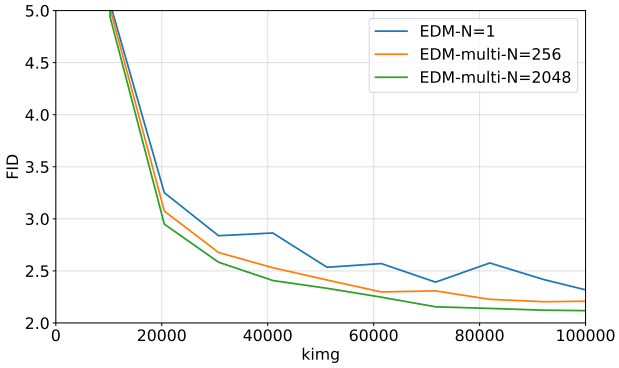

(b) EDM. Larger N alters both convergence speed and final FID.

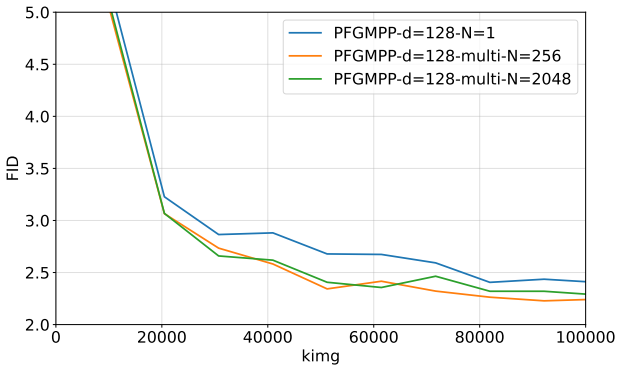

(c) PFGM++ ($d = 128$). Larger N improves convergence and final FID.

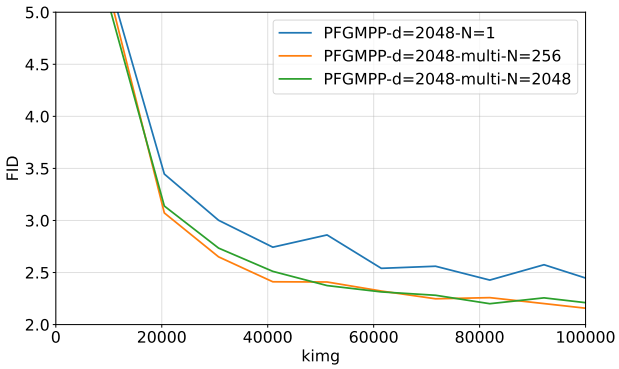

(d) PFGM++ ($d = 2048$). Performance depends noticeably on N, with larger target sets yielding stronger results.

Figure 5. Training convergence for CIFAR-10 dataset (FID vs. kimg) for different generative methods and multi-sample target sizes N. One kimg = one thousand images processed.

In Figure 6, we present the training convergence in terms of the FID metric for the FFHQ-64 dataset. Final checkpoint FID values are reported in Table 6. Compared to the CIFAR-10 dataset with 32×32 images, the improvement in FID for FFHQ at 64×64 resolution is noticeably smaller.

This behavior can be explained through the Gini coefficient computed from the weights in Eq. (31):

$$G = 1 - \sum_{i=1}^{N} w_i^2. \tag{35}$$

Lower Gini coefficient values indicate that one weight dominates the others. In our setting, this means that the multi-sample target effectively approaches a single-sample target. Table 7 shows that increasing image resolution leads to lower Gini coefficient values, suggesting that for higher-resolution data the multi-sample target becomes increasingly similar to the single-sample target.

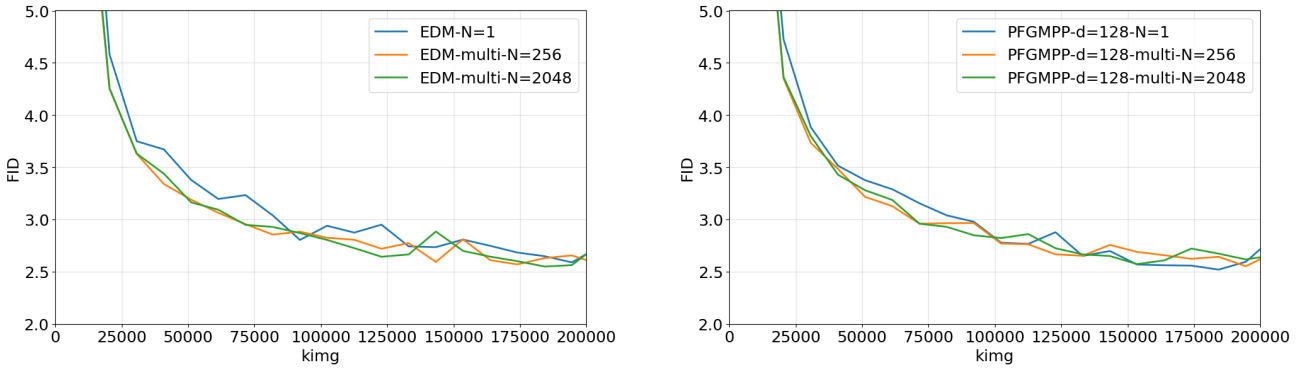

(a) FID for EDM with single target and multi-sample target with $N = 256$ and $N = 2048$ for FFHQ 64x64 dataset

(b) FID for PFGM++ with single target and multi-sample target with $N = 256$ and $N = 2048$ for FFHQ 64x64 dataset

Figure 6. FID for different kimg on training for FFHQ 64x64 dataset

Table 6. Training convergence for FFHQ 64x64 dataset (FID vs. kimg) for different generative methods and multi-sample target sizes N. One kimg = one thousand images processed. Multi-Sample target gives statistically meaningful improvement.

N	**FID** (↓)	
	EDM	**PFGM++** ($d = 128$)
1	2.62 ± 0.04	2.69 ± 0.02
256	$\mathbf{2.51 \pm 0.05}$	$\mathbf{2.56 \pm 0.03}$
2048	2.63 ± 0.04	2.61 ± 0.03

Table 7. Gini coefficient (35) for resized CIFAR resolutions and different multi-sample target sizes N. Higher image dimensionality leads to lower Gini values, indicating that the multi-sample target becomes closer to the single-sample target.

N	**Resized CIFAR Resolution**			
	8x8	**16x16**	**32x32**	**64x64**
128	0.214	0.098	0.005	0.001
256	0.211	0.112	0.005	0.001
512	0.214	0.125	0.006	0.001
1024	0.216	0.139	0.007	0.001
2048	0.215	0.154	0.009	0.002

Limitations

Our main contribution is theoretical. The empirical results are intended to validate the duality and illustrate estimator behavior, not to establish a new state-of-the-art generative model.