Unlocking the Duality between Flow and Field Matching

Daniil Shlenskii
AXXX, Moscow, Russia
Applied AI Institute, Moscow, Russia
daniil.shlenskii@gmail.com

Alexander Varlamov
Applied AI Institute, Moscow, Russia

Nazar Buzun
AXXX, Moscow, Russia
Innopolis University, Innopolis, Russia

Alexander Korotin
Applied AI Institute, Moscow, Russia
AXXX, Moscow, Russia

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 give explicit constructions in both directions and show that CFM and forward-only IFM induce the same family of data-space dynamics. General IFM, however, is strictly more expressive: already in EFM, its sampling rule yields transport maps that reverse orientation with positive probability, something no CFM flow can do. Our findings suggest using the probabilistic and geometric views jointly when developing generative models, and exploring field-based transport beyond the forward-only regime.
Code: https://github.com/daniil-shlenskii/flow-field-matching-duality.

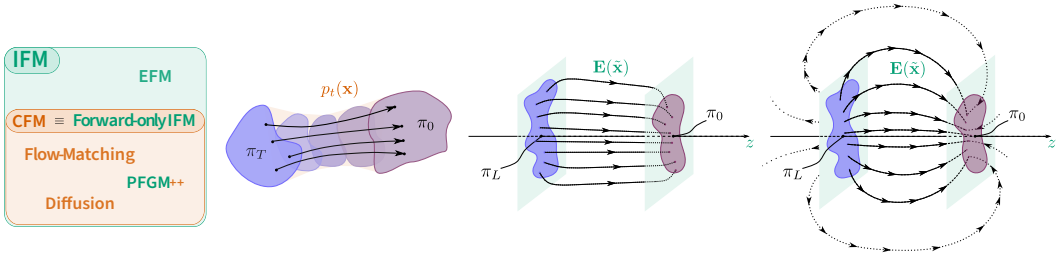

(a) Flow-field taxonomy (b) CFM (c) Forward-only IFM ≡ CFM (d) EFM ∈ IFM \ CFM

Figure 1: **Overview of the flow–field duality.** We establish that CFM and forward-only IFM induce the same data-space generative dynamics (§3): CFM $\equiv$ forward-only IFM $\subsetneq$ 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 more expressive: it includes EFM (d, §2.2), which admits returning trajectories whose transport maps reverse orientation and therefore cannot be realized by any CFM flow (§4).

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

1 Introduction

Modern continuous-time generative modeling is largely organized around flow-based descriptions of transport. Diffusion [1, 2] and flow-based [3–5] generative models can both be unified under the Conditional Flow Matching [6, 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 [7, 8], field-based generative models, including Poisson Flow Generative Models [9, PFGM], PFGM++ [10], and Electrostatic Field Matching [11, 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 [12, 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 establish the equivalence between CFM and forward-only IFM, a natural subclass of IFM (§2.2), through explicit constructions in both directions (§3.2, §3.3). These constructions preserve the induced data-space dynamics, providing probabilistic and geometric descriptions of the same generative process.
- **A strict separation between CFM and general IFM.** We show that general IFM is strictly more expressive than CFM at the level of data-space transport maps (§4): while CFM transport maps never reverse orientation, general IFM admits maps that do. We also show that EFM is a particular instance of IFM realizing this property.

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 orientation reversal as a concrete mechanism by which general IFM goes beyond CFM, motivating field-based generative modeling outside the forward-only regime as a new research direction.

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 ODE-based dynamics, 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 [13]. 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 [6]. The one-sided setting is common in noise-to-data generation [14], 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 [3, 15].

Definition (Conditional Flow). A *Conditional Flow* [6] 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 $t \in (0, T)$

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

where $\nabla_{\mathbf{x}}\cdot$ denotes the divergence operator with respect to $\mathbf{x}$, together with the *boundary conditions*

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

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**

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

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

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

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

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

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 [16, §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 [16, §3.2], also known as Flow Matching [4], 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 introduced in [4]. The same choice is also closely related to the Gaussian flow-matching formulation of diffusion models [13].

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 connect the two plates and, traversed from $z = L$ to $z = 0$, define the transport from the source π_L to the target 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 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* [12] if it satisfies the following three properties.

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

$$\int_{\partial M} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}}) \cdot \mathbf{dS} \;=\; \Phi_0 \tag{6}$$

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

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, all **field lines**, except for a family carrying zero flux through every transversal hypersurface [1], connect the two particles:

$$d\tilde{\mathbf{x}}(\tau)/d\tau = \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}}(\tau)), \quad \lim_{\tau \downarrow \tau_s} \tilde{\mathbf{x}}(\tau) = \tilde{\mathbf{x}}_0, \quad \lim_{\tau \uparrow \tau_f} \tilde{\mathbf{x}}(\tau) = \tilde{\mathbf{x}}_L, \tag{8}$$

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

Remark. The parameter τ in (8) 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 $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})$.

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 [12], one can transport samples from π_L to π_0 following the field lines of $\mathbf{E}$:

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

In general, simulating (10) poses practical difficulties: **(a)** 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 **(b)** field lines may leave the capacitor region $0 \le z \le L$ and cross the plates several times, which requires the stochastic starting and stopping rules of Manukhov et al. [12, Appendix A.2]. *For a forward-only interaction field,*

[1] A hypersurface $S \subset \mathbb{R}^{D+1}$ is called *transversal* if $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L} \cdot \mathbf{n} > 0$ along it, where $\mathbf{n}$ is a unit normal to S.

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:

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

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

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

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

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

Motivated by this observation, Manukhov et al. [12] use the following training objective:

$$\mathbb{E}_{\tilde{\mathbf{x}}\sim p_{\text{vol}}} \|f_\theta(\tilde{\mathbf{x}}) - \mathbf{E}(\tilde{\mathbf{x}})/\|\mathbf{E}(\tilde{\mathbf{x}})\|_2\|_2^2\,, \quad \mathbf{E}(\tilde{\mathbf{x}}) = (\mathbf{E}(\tilde{\mathbf{x}})_{\mathbf{x}}, \mathbf{E}(\tilde{\mathbf{x}})_z) \tag{13}$$

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 [9, 11], 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 defined above. Moreover, its field lines may move backward along the z-axis and leave the capacitor region $0 \le z \le L$; the EFM field is therefore not forward-only, and sampling uses the general IFM stochastic rule.

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 [12, 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.2 for convenience.

2.3 CFM vs. IFM

Previously, we reviewed the CFM and IFM frameworks separately. Although their differences may already be apparent, we provide **Table 1** in Appendix A.1 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** (13) and learns it at points drawn from a 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:

$$\mathbf{E}^{\mathbf{x}_0}(\tilde{\mathbf{x}}) \propto (\tilde{\mathbf{x}} - \tilde{\mathbf{x}}_0)/\|\tilde{\mathbf{x}} - \tilde{\mathbf{x}}_0\|_2^{D+1}, \tag{14}$$

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 an instance of the two-sided interaction field defined above: its conditional field is one-sided, constructed from a single target point $\mathbf{x}_0$. In §3.1 we introduce a one-sided notion of an interaction field, of which PFGM is an instance.

Flow-based interpretation. Subsequent work [10, 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

$$\mathbb{E}_{z,\mathbf{x}_0,\mathbf{x}_z} \, \|f_\theta(\mathbf{x}_z, z) - (\mathbf{x}_z - \mathbf{x}_0)/z\|_2^2, \quad p_z^{\mathbf{x}_0}(\mathbf{x}) \propto \left(\|\mathbf{x} - \mathbf{x}_0\|_2^2 + z^2\right)^{-\frac{D+1}{2}}, \tag{15}$$

where $z \sim \mathcal{U}[0, L]$ and $\mathbf{x}_z \sim p_z^{\mathbf{x}_0}$. The similarity between (15) and the CFM objective (5) suggests a flow-based interpretation of PFGM: $p_z^{\mathbf{x}_0}$ in (15) 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 one generative transport can be described either as a *flow* or as *field lines*. This motivates three questions:

(1) PFGM is one-sided: can IFM be given a one-sided formulation, as CFM has? (§3.1)

(2) Is this dual description specific to PFGM, or is there a CFM–IFM duality? (§3.2, §3.3)

(3) Can general IFM realize data-space transport beyond this duality? (§4)

3 Flow–field duality: CFM ≡ forward-only IFM

In this section, we take up the first two questions raised in §2.4 and 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. Together, the two directions establish the equivalence; the third question, what lies outside the forward-only class, is taken up in §4. All proofs are provided in Appendix B.

Before stating the duality, we make a simple observation. In forward-only fields, the spatial coordinate z (the distance from the target plate) plays the same role as the time variable t in conditional flows. Likewise, the plate separation L corresponds to the time horizon T. With this correspondence in mind, we will use the (t, T)-based notation for conditional flows and interaction fields alike 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 a *one-sided Interaction Field* if it satisfies the following three properties. First, the field has a fixed total flux (6) through any closed surface that encloses $\tilde{\mathbf{x}}_0$ and lies in the region $\{t < T\}$. Second, the field is divergence-free: $\nabla \cdot \mathbf{E}^{\mathbf{x}_0}(\tilde{\mathbf{x}}) = 0$ for all $\tilde{\mathbf{x}}$ with $t < T$, $\tilde{\mathbf{x}} \neq \tilde{\mathbf{x}}_0$. Third, all its field lines, except for a family carrying zero flux through every transversal hypersurface (§2.2), start at the particle $\tilde{\mathbf{x}}_0 = (\mathbf{x}_0, 0)$: $d\tilde{\mathbf{x}}(\tau)/d\tau = \mathbf{E}^{\mathbf{x}_0}(\tilde{\mathbf{x}}(\tau))$, $\lim_{\tau \downarrow \tau_s} \tilde{\mathbf{x}}(\tau) = \tilde{\mathbf{x}}_0$, where τ_s denotes the starting value of the field-line parameter, and reach the plate $\{t = T\}$: there exists $\tau_f > \tau_s$ such that $\tilde{\mathbf{x}}(\tau_f) \in \mathbb{R}^D \times \{T\}$.

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 Flows 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*

$$\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t) := (\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}), p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})). \tag{16}$$

Then, under the assumptions specified in Appendix B.2:

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*
$$\mathbf{E}(\mathbf{x}, t) = \big(\mathbf{v}_t(\mathbf{x})\, p_t(\mathbf{x}),\ p_t(\mathbf{x})\big). \tag{17}$$
3. *The CFM generative dynamics* (4) *coincide with the dynamics induced by field lines* (10)*, reparameterized by the distance from the plate* (12)*:*
$$d\mathbf{x}/dt = \mathbf{v}_t(\mathbf{x}) = \mathbf{E}(\mathbf{x}, t)_{\mathbf{x}}/\mathbf{E}(\mathbf{x}, t)_t. \tag{18}$$

Our Theorem 3.1 shows that every 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 B.1, Table 2.

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

$$\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t)_{\mathbf{x}}/\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t)_t, \quad p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t)_t. \tag{19}$$

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

$$\Phi_t^{\mathbf{x}_0,\mathbf{x}_T} := \int_{\mathbb{R}^D} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t)_t\, d\mathbf{x} \tag{20}$$

need not equal 1, and may, a priori, depend on $\mathbf{x}_0, \mathbf{x}_T$ and t. Nevertheless, it is in fact independent of both, as follows from the proposition below.

Proposition 3.2. (Flux Conservation across Slices). *Under the assumptions specified in Appendix B.4, the scalar $\Phi_t^{\mathbf{x}_0,\mathbf{x}_T}$ in* (20) *is the flux through the slice $\{t = \text{const}\}$ at distance t from the target plate and coincides with the total flux Φ_0* (6) *of the forward-only Interaction Field.*

Since Φ_0 is a positive constant, dividing by it rescales the field without changing its lines, and hence the induced dynamics. We now establish the converse 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*

$$\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_{\mathbf{x}}/\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t, \quad p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t/\Phi_0. \tag{21}$$

Then, under the assumptions specified in Appendix B.4:

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*

$$\mathbf{v}_t(\mathbf{x}) = \mathbf{E}(\mathbf{x},t)_{\mathbf{x}}/\mathbf{E}(\mathbf{x},t)_t, \quad p_t(\mathbf{x}) = \mathbf{E}(\mathbf{x},t)_t/\Phi_0. \tag{22}$$

3. *Consequently, as in Theorem 3.1, the CFM generative dynamics* (4) *coincide with the dynamics induced by field lines* (10)*, reparameterized by the distance from the plate* (12)*.*

Theorem 3.3 shows that every 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 noted 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 B.1, Table 2. 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 the same table.

4 Beyond forward-only IFM

The constructions in §3 establish an equivalence between CFM and forward-only IFM at the conditional level. Since they commute with marginalization and superposition, this equivalence extends to the induced global data-space dynamics. This leaves open whether general IFM can realize data-space transport beyond this equivalence. Backward-oriented field lines alone do not establish such a separation: different geometries in the augmented space may still induce the same data-space transport. All proofs are provided in Appendix C.

To prove a strict separation, we therefore identify an obstruction at the level of data-space transport maps. We approach this question through orientation: a differentiable map $G : \mathbb{R}^D \to \mathbb{R}^D$ preserves (reverses) orientation on an open set if $\det \nabla G > 0$ (< 0) there. In one dimension, this means preserving (reversing) the order of nearby points; for example, $G(a) = -a$ is orientation-reversing.

Let $G_{s,t}$ denote the flow map of the global velocity $\mathbf{v}_t$ from time s to t, i.e., $\partial_t G_{s,t}(a) = \mathbf{v}_t(G_{s,t}(a))$ with $G_{s,s}(a) = a$. Since $\mathbf{v}_t$ may be singular at an endpoint carrying a compactly supported density [17, §3.5], we only assume it to be smooth on $(0,T)$ and take the CFM transport map $G_{T,0}$ to be the locally uniform limit of $G_{s,t}$ as $s \uparrow T$, $t \downarrow 0$, wherever it exists. Each $G_{s,t}$ with $0 < t < s < T$ is smooth and injective [17, Thm. 1], with $\det \nabla G_{s,t}(a) = \exp\big(\int_s^t (\nabla \cdot \mathbf{v}_\tau)(G_{s,\tau}(a))\, d\tau\big) > 0$ by Liouville's formula. A locally uniform limit of such orientation-preserving maps cannot reverse orientation: $\det \nabla G_{T,0} \geq 0$ wherever $G_{T,0}$ is continuously differentiable (a standard consequence of degree theory [18, Def. 2.1 and Thm. 3.1]).

General IFM admits a different behavior. A generative trajectory starting at $t = T$ may reach the target plate $t = 0$, continue into $t < 0$, and then return to $t = 0$. As we show below, the resulting transport map reverses orientation.

Proposition 4.1 (Returning trajectories reverse orientation)**.** *Let* $\mathbf{E}$ *be a global IFM field and let* $U \subset \mathbb{R}^D$ *be a nonempty open set. Suppose the generative trajectories issued from* $U \times \{T\}$ *enter the region between the plates, cross the target plate, and return to it from* $t < 0$*, and let* $F : U \to \mathbb{R}^D$ *be the resulting transport map. Then, under regularity and transversality assumptions specified in Appendix C.1,* F *reverses orientation.*

The next proposition shows that EFM realizes this mechanism: its sampling rule selects such a returning trajectory with positive probability.

Proposition 4.2 (EFM admits orientation-reversing trajectories)**.** *For every* $D \geq 1$*, there exist smooth, compactly supported source and target probability densities and a nonempty open set* $U \subset \mathbb{R}^D$ *with* $\pi_T(U) > 0$ *such that, under the IFM sampling rule [12, Appendix A.2] in the conventions of Appendix C.3, a trajectory initialized at any* $a \in U$ *crosses the target plate, enters* $t < 0$*, returns, and terminates at its second crossing with probability one. The resulting transport map* $F : U \to \mathbb{R}^D$ *reverses orientation:* $\det \nabla F(a) < 0$ *for every* $a \in U$*.*

Together with the orientation property of CFM, Proposition 4.2 yields the desired separation.

Corollary 4.3 (Strict separation)**.** *The EFM transport map* F *of Proposition 4.2 is realized by no CFM flow, for any source and target densities:* $G_{T,0}$ *and* F *do not coincide on any nonempty open subset of* U*. Combined with §3, general IFM is therefore strictly more expressive than CFM at the level of data-space transport maps.*

Indeed, if $G_{T,0} = F$ on a nonempty open $V \subset U$, then $G_{T,0}$ is continuously differentiable on V with $\det \nabla G_{T,0} = \det \nabla F < 0$, contradicting the orientation property of CFM.

5 Duality Takeaways

We now collect the main implications of the flow–field duality. The section is organized around three takeaways: for each one we state the implication and briefly discuss why it matters. Each is supported empirically on a controlled setting, described together with the results in Appendix D. Additional discussions are provided in Appendix E.

Takeaway 1 (Expressiveness). At the level of induced data-space transport maps,

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

Theorems 3.1 and 3.3 show that CFM and forward-only IFM induce 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 trajectories that cross the target plate and return: Proposition 4.1 shows that their transport maps reverse orientation, whereas no CFM transport map does. This already occurs in EFM, with positive probability (Proposition 4.2). This is the concrete property that puts general IFM outside CFM, and the reason it is worth developing on its own.

Appendix D.2 verifies the equivalence of §3 directly: we train a CFM drift and its dual interaction field independently, and compare the data-space dynamics they induce. Built from the same conditional flow, the two agree to a mean squared error of order 10^{-7}; dynamics built from different conditional flows differ by 4.9×10^{-1}, six orders of magnitude more. The equivalence therefore survives estimation: it is not only an identity between exact objects, but something two separately trained models arrive at.

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:

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

Theorem 3.3 shows that the geometric IFM representation has an equivalent probabilistic view. Namely, for the forward-only fields the information about the induced probabilistic path is carried by the last component of the global field.

Appendix D.3 illustrates this correspondence empirically. We compare particle distributions generated by the trained IFM model with heatmaps of the last component of the global field, estimated by Monte Carlo averaging over sampled endpoint pairs. Their qualitative agreement (Figure 3) supports 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

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

Training an interaction field requires a distribution p_{vol} covering the region where the field is defined (13), and IFM offers no principled choice: it is picked heuristically, one of the differences from CFM noted in §2.3.

Substituting the duality (21) into the CFM objective (5) gives

$$\mathbb{E}_{t, \mathbf{x}_0, \mathbf{x}_T,\, \mathbf{x} \sim p_t^{\mathbf{x}_0, \mathbf{x}_T}} \left\| \mathbf{v}_\theta(\mathbf{x}, t) - \mathbf{E}^{\mathbf{x}_0, \mathbf{x}_T}(\mathbf{x}, t)_{\mathbf{x}} / \mathbf{E}^{\mathbf{x}_0, \mathbf{x}_T}(\mathbf{x}, t)_t \right\|_2^2 . \tag{23}$$

This has two consequences. First, the dynamics induced by a forward-only interaction field can be learned by single-sample regression: the target in (23) is built from one endpoint pair, whereas the field target in (13) has to be estimated from a batch. Second, the objective arrives with its own sampling distribution $p_t^{\mathbf{x}_0, \mathbf{x}_T} = \mathbf{E}^{\mathbf{x}_0, \mathbf{x}_T}(\mathbf{x}, t)_t / \Phi_0$, which is exactly the principled p_{vol} that IFM was missing. It can be used inside IFM as well, replacing the heuristic volume distribution in the normalized-field objective:

$$\mathbb{E}_{t, \mathbf{x}_0, \mathbf{x}_T,\, \mathbf{x} \sim p_t^{\mathbf{x}_0, \mathbf{x}_T}} \left\| f_\theta(\mathbf{x}, t) - \mathbf{E}(\mathbf{x}, t) / \|\mathbf{E}(\mathbf{x}, t)\|_2 \right\|_2^2 . \tag{24}$$

Appendix D.4 checks the second consequence: we fix the normalized-field objective (13) and the target field, and vary only p_{vol}, comparing the field's own conditional path, which turns the objective into (24), with the path induced by a different field. Its own path more than halves the final W_2 for both fields we test ($0.58 \to 0.25$ and $0.81 \to 0.38$). The choice of p_{vol} is therefore not a harmless implementation detail: it costs accuracy when it does not match the target field, and the duality supplies the matching one without an additional heuristic choice.

Conclusion. In this work we reveal the relation between two frameworks, CFM and IFM, that are built on different principles. We show that on the forward-only subclass they are one theory, constructively and in both directions. This duality removes the reason to study the two frameworks apart in that regime, and its field-to-flow direction gives IFM two things it lacked: an explicit marginal path and a training distribution read off the field itself. The inclusion, however, is strict: general IFM is not a reformulation of CFM but a genuinely larger class, so models outside the forward-only subclass can do things no conditional flow can, which makes general IFM a direction worth developing on its own. The separating property is concrete, orientation reversal, and EFM already realizes it, which completes the taxonomy in Figure 1.

Acknowledgments

The work was supported by the grant for research centers in the field of AI provided by the Ministry of Economic Development of the Russian Federation in accordance with the agreement 000000C313925P4F0002 and the agreement No. 139-10-2025-033.

References

[1] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. pmlr, 2015.

[2] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.

[3] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.

[4] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

[5] Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022.

[6] Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *arXiv preprint arXiv:2302.00482*, 2023.

[7] Ziming Liu, Di Luo, Yilun Xu, Tommi Jaakkola, and Max Tegmark. Genphys: From physical processes to generative models. *arXiv preprint arXiv:2304.02637*, 2023.

[8] Daniil Shlenskii and Alexander Korotin. Overclocking electrostatic generative models. In *Forty-third International Conference on Machine Learning*, 2026. URL `https://openreview.net/forum?id=flo449mncA`.

[9] Yilun Xu, Ziming Liu, Max Tegmark, and Tommi Jaakkola. Poisson flow generative models. *Advances in Neural Information Processing Systems*, 35:16782–16795, 2022.

[10] Yilun Xu, Ziming Liu, Yonglong Tian, Shangyuan Tong, Max Tegmark, and Tommi Jaakkola. Pfgm++: Unlocking the potential of physics-inspired generative models. In *International Conference on Machine Learning*, pages 38566–38591. PMLR, 2023.

[11] Alexander Kolesov, Stepan Manukhov, Vladimir V. Palyulin, and Alexander Korotin. Field matching: an electrostatic paradigm to generate and transfer data. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 31202–31222. PMLR, 2025. URL `https://proceedings.mlr.press/v267/kolesov25a.html`.

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

[13] Ruiqi Gao, Emiel Hoogeboom, Jonathan Heek, Valentin De Bortoli, Kevin Patrick Murphy, and Tim Salimans. Diffusion models and gaussian flow matching: Two sides of the same coin. In *The Fourth Blogpost Track at ICLR 2025*, 2025.

[14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

[15] Michael S Albergo, Mark Goldstein, Nicholas M Boffi, Rajesh Ranganath, and Eric Vanden-Eijnden. Stochastic interpolants with data-dependent couplings. *arXiv preprint arXiv:2310.03725*, 2023.

[16] Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.

[17] Yaron Lipman, Marton Havasi, Peter Holderrieth, Neta Shaul, Matt Le, Brian Karrer, Ricky TQ Chen, David Lopez-Paz, Heli Ben-Hamu, and Itai Gat. Flow matching guide and code. *arXiv preprint arXiv:2412.06264*, 2024.

[18] Klaus Deimling. *Nonlinear Functional Analysis*. Springer-Verlag, Berlin, Heidelberg, 1985.

[19] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35: 26565–26577, 2022.

[20] Vladimir I Bogachev, Michael Röckner, and Stanislav V Shaposhnikov. 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.

[21] Gleb Ryzhakov, Svetlana Pavlova, Egor Sevriugov, and Ivan Oseledets. Explicit flow matching: On the theory of flow matching algorithms with applications. *arXiv preprint arXiv:2402.03232*, 2024.

[22] Quentin Bertrand, Anne Gagneux, Mathurin Massias, and Rémi Emonet. On the closed-form of flow matching: Generalization does not arise from target stochasticity. *arXiv preprint arXiv:2506.03719*, 2025.

[23] Yilun Xu, Shangyuan Tong, and Tommi Jaakkola. Stable target field for reduced variance score estimation in diffusion models. In *The Eleventh International Conference on Learning Representations*, 2023.

[24] Alex Krizhevsky. 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`.

[25] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4401–4410, June 2019.

[26] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium, 2018. URL `https://arxiv.org/abs/1706.08500`.

A Additional material for Section 2

A.1 CFM vs. IFM: side-by-side comparison

Table 1 lists the main terms of the two frameworks side by side, as discussed in §2.3.

Table 1: **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** (8) $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})$ 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**: $\frac{d\mathbf{x}}{dt} = \mathbf{v}_t(\mathbf{x}), \quad \mathbf{x}_T \sim \pi_T$	Field-line transport **in the extended space**: $\frac{d\tilde{\mathbf{x}}}{d\tau} = \mathbf{E}(\tilde{\mathbf{x}}), \quad \tilde{\mathbf{x}} \sim \pi_L \times \delta_{z=L}$ **Forward-only** fields dynamics **in the data space**: $\frac{d\mathbf{x}}{dz} = \frac{\mathbf{E}(\tilde{\mathbf{x}})_{\mathbf{x}}}{\mathbf{E}(\tilde{\mathbf{x}})_z}, \quad \mathbf{x}_L \sim \pi_L$
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: $\mathbb{E}_{\tilde{\mathbf{x}}\sim p_{\mathrm{vol}}} \|f_\theta(\tilde{\mathbf{x}}) - \mathbf{E}(\tilde{\mathbf{x}})\|_2^2$ **Normalized field** in the forward-only case: $\mathbb{E}_{\tilde{\mathbf{x}}\sim p_{\mathrm{vol}}} \left\| f_\theta(\tilde{\mathbf{x}}) - \frac{\mathbf{E}(\tilde{\mathbf{x}})}{\|\mathbf{E}(\tilde{\mathbf{x}})\|_2} \right\|_2^2$
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: $\tilde{\mathbf{x}} \sim p_{\mathrm{vol}}$

A.2 Original IFM field realization

For completeness, we restate the forward-only Interaction Field realization proposed in Manukhov et al. [12, Appendix A.4]. They construct $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}$ in two steps. For a pair on a common axis, $\tilde{\mathbf{x}}_0 = (\mathbf{0}, 0)$ and $\tilde{\mathbf{x}}_L = (\mathbf{0}, L)$, split an arbitrary point as $\tilde{\mathbf{x}} = \tilde{\mathbf{x}}_\perp + z\,\mathbf{e}_z$, where $\mathbf{e}_z$ is the unit vector along the z-axis and $\tilde{\mathbf{x}}_\perp$ is the component of $\tilde{\mathbf{x}}$ orthogonal to it, and write $x_\perp := \|\tilde{\mathbf{x}}_\perp\|_2$ and $\mathbf{e}_\perp := \tilde{\mathbf{x}}_\perp / x_\perp$. The field is then given by a unit direction $\mathbf{n}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})$ and a magnitude $\|\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}})\|_2$:

$$\mathbf{n}^{\mathbf{x}_0,\mathbf{x}_L}(\tilde{\mathbf{x}}) := \cos\alpha(\tilde{\mathbf{x}})\,\mathbf{e}_z + \sin\alpha(\tilde{\mathbf{x}})\,\mathbf{e}_\perp, \tag{25}$$

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

$$\sigma(z) := \sin\left(\frac{\pi z}{L}\right), \qquad \alpha(\tilde{\mathbf{x}}) := \arctan\left(\frac{\pi}{L}\, x_\perp \cot\left(\frac{\pi z}{L}\right)\right), \tag{27}$$

where $\sigma(z)$ is the width of the tube at height z (for $d = L/2$ and $\sigma_0 = 1$ in their notation) and the angle $\alpha(\tilde{\mathbf{x}})$, measured from the z-axis, is chosen so that the ratio $x_\perp/\sigma(z)$ remains constant along each field line. For a pair in general position, $\tilde{\mathbf{x}}_0 = (\mathbf{x}_0, 0)$ and $\tilde{\mathbf{x}}_L = (\mathbf{x}_L, L)$, the field is obtained from the above by a shift parallel to the plates, carrying the axis onto the segment joining $\tilde{\mathbf{x}}_0$ and $\tilde{\mathbf{x}}_L$ and leaving the height z unchanged.

B Additional material for Section 3

This section collects the dual constructions referenced in §3 together with the assumptions and proofs for its results.

B.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 2.

Table 2: **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 [4]	$\mathbf{v}_t^{\mathbf{x}_0}(\mathbf{x}) = \dfrac{\mathbf{x}-\mathbf{x}_0}{t}$ $p_t^{\mathbf{x}_0}(\mathbf{x}) = \mathcal{N}(\mathbf{x} \mid (1-t)\mathbf{x}_0,\ t^2 I)$	$\mathbf{E}^{\mathbf{x}_0}(\mathbf{x},t)_{\mathbf{x}} = \mathbf{v}_t^{\mathbf{x}_0}(\mathbf{x})\, p_t^{\mathbf{x}_0}(\mathbf{x})$ $\mathbf{E}^{\mathbf{x}_0}(\mathbf{x},t)_t = p_t^{\mathbf{x}_0}(\mathbf{x})$
VE-diffusion / EDM [2, 19]	$\mathbf{v}_t^{\mathbf{x}_0}(\mathbf{x}) = \dfrac{\dot\sigma(t)}{\sigma(t)}(\mathbf{x}-\mathbf{x}_0)$ $p_t^{\mathbf{x}_0}(\mathbf{x}) = \mathcal{N}(\mathbf{x} \mid \mathbf{x}_0,\ \sigma(t)^2 I)$	$\mathbf{E}^{\mathbf{x}_0}(\mathbf{x},t)_{\mathbf{x}} = \mathbf{v}_t^{\mathbf{x}_0}(\mathbf{x})\, p_t^{\mathbf{x}_0}(\mathbf{x})$ $\mathbf{E}^{\mathbf{x}_0}(\mathbf{x},t)_t = p_t^{\mathbf{x}_0}(\mathbf{x})$
Linear Interpolant [16]	$\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \dfrac{\dot\sigma(t)}{\sigma(t)}(\mathbf{x}-\mathcal{I}^{\mathbf{x}_0,\mathbf{x}_T}(t)) + \dot{\mathcal{I}}^{\mathbf{x}_0,\mathbf{x}_T}(t)$ $p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \mathcal{N}(x \mid \mathcal{I}(t,\mathbf{x}_0,\mathbf{x}_T), \sigma(t)^2 I)$ $\mathcal{I}^{\mathbf{x}_0,\mathbf{x}_T}(t) = \mathbf{x}_0\cdot(1-t/T) + \mathbf{x}_T\cdot t/T$ $\sigma(t) = \sqrt{2\cdot t/T\cdot(1-t/T)}$	$\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_{\mathbf{x}} = \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$ $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t = p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$
PFGM [9]	$\mathbf{v}_t^{\mathbf{x}_0}(\mathbf{x}) = \dfrac{\mathbf{x}-\mathbf{x}_0}{t}$ $p_t^{\mathbf{x}_0}(\mathbf{x}) \propto t/\left(\|\mathbf{x}-\mathbf{x}_0\|_2^2 + t^2\right)^{\frac{D+1}{2}}$	$\mathbf{E}^{\mathbf{x}_0}(\tilde{\mathbf{x}}) = \dfrac{\tilde{\mathbf{x}}-\tilde{\mathbf{x}}_0}{\|\tilde{\mathbf{x}}-\tilde{\mathbf{x}}_0\|_2^{D+1}}$
PFGM++ ($d \geq 1$) [10]	$\mathbf{v}_t^{\mathbf{x}_0}(\mathbf{x}) = \dfrac{\mathbf{x}-\mathbf{x}_0}{t}$ $p_t^{\mathbf{x}_0}(\mathbf{x}) \propto t^d/\left(\|\mathbf{x}-\mathbf{x}_0\|_2^2 + t^2\right)^{\frac{D+d}{2}}$	$\mathbf{E}^{\mathbf{x}_0}(\mathbf{x},t)_{\mathbf{x}} = \mathbf{v}_t^{\mathbf{x}_0}(\mathbf{x})\, p_t^{\mathbf{x}_0}(\mathbf{x})$ $\mathbf{E}^{\mathbf{x}_0}(\mathbf{x},t)_t = p_t^{\mathbf{x}_0}(\mathbf{x})$
IFM [12]	$\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \dfrac{\dot\sigma(t)}{\sigma(t)}(\mathbf{x}-\mathcal{I}^{\mathbf{x}_0,\mathbf{x}_T}(t)) + \dot{\mathcal{I}}^{\mathbf{x}_0,\mathbf{x}_T}(t)$ $p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \mathcal{N}(x \mid \mathcal{I}(t,\mathbf{x}_0,\mathbf{x}_T), \sigma(t)^2 I)$ $\mathcal{I}^{\mathbf{x}_0,\mathbf{x}_T}(t) = \mathbf{x}_0\cdot(1-t/T) + \mathbf{x}_T\cdot t/T$ $\sigma(t) = \sin(\pi t/T)$	$\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_{\mathbf{x}} = \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$ $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t = p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$ OR originally defined as in §2.2
EFM [11]	$\times$	$\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\tilde{\mathbf{x}}) = \dfrac{\tilde{\mathbf{x}}-\tilde{\mathbf{x}}_0}{\|\tilde{\mathbf{x}}-\tilde{\mathbf{x}}_0\|_2^{D+1}} - \dfrac{\tilde{\mathbf{x}}-\tilde{\mathbf{x}}_T}{\|\tilde{\mathbf{x}}-\tilde{\mathbf{x}}_T\|_2^{D+1}}$

B.2 Assumptions for Theorem 3.1

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

1. $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} = (\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T} p_t^{\mathbf{x}_0,\mathbf{x}_T},\ p_t^{\mathbf{x}_0,\mathbf{x}_T}) \in C^1((0,T)\times\mathbb{R}^D)$, and $p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$ 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; in particular $\int_0^T\int_{\mathbb{R}^D} \|\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\|\, p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\, d\mathbf{x}\, dt < \infty$, which is the integrability hypothesis of the superposition principle.

Boundary convention. In both directions of the duality, pointwise identities are understood on the open slab $\mathbb{R}^D \times (0,T)$. After extension by zero, boundary charges and fluxes are understood in the distributional and weak normal-trace senses, respectively; no smooth extension across the plates is required.

B.3 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, $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t = p_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$, which is strictly positive on $(0,T)\times\mathbb{R}^D$ by Assumption 1. Outside the slab the field is extended by zero, so $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)_t > 0$ at every point where $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}$ is nonzero, and $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}$ is a forward-only field.

   Field lines connect the two particles. The connecting property (8) follows from the superposition principle together with the boundary conditions (2).

   By the superposition principle for the continuity equation [20], applied to the pair $(\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}, p_t^{\mathbf{x}_0,\mathbf{x}_T})$ solving (1), there exists a probability measure η on the space of continuous curves $C([0,T];\mathbb{R}^D)$, concentrated on absolutely continuous curves $\mathbf{x}(\cdot)$ that solve

   $$\frac{d\mathbf{x}}{dt} = \mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}(t)) \qquad \text{for a.e. } t\in(0,T),$$

   and such that $(e_t)_\#\eta = p_t^{\mathbf{x}_0,\mathbf{x}_T}$ for every $t\in[0,T]$, where $e_t\big(\mathbf{x}(\cdot)\big) := \mathbf{x}(t)$ denotes the evaluation map.

   The boundary conditions (2) read $(e_0)_\#\eta = \delta_{\mathbf{x}_0}$ and $(e_T)_\#\eta = \delta_{\mathbf{x}_T}$. A push-forward equal to a Dirac mass forces the corresponding evaluation to be η-almost surely constant. Let C be the set of absolutely continuous solutions of the above ODE with $\mathbf{x}(0)=\mathbf{x}_0$ and $\mathbf{x}(T)=\mathbf{x}_T$. Then $\eta(C)=1$. Thus η-almost every curve of the representation runs from $\mathbf{x}_0$ to $\mathbf{x}_T$; in particular, at least one such curve exists. Such a curve need not be unique, and Definition (8) does not require uniqueness: it only asks the field lines to *connect* the two particles, up to a family carrying zero flux. In typical constructions the connecting family is in fact infinite (Remark B.1).

   Fix any curve $\mathbf{x}(\cdot)\in C$ and put $\tilde{\mathbf{x}}(t) := (\mathbf{x}(t),t)$. It 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

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

   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

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

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

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

   The reparameterization is well defined since $p_t^{\mathbf{x}_0,\mathbf{x}_T}$ is strictly positive and C^1 (Assumption 1); the resulting parameter range (τ_s,τ_f) may be unbounded, which Definition (8) allows, as it constrains only the limits at its endpoints. Therefore $\tilde{\mathbf{x}}(\tau)$ is a field line of $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}$ connecting $\tilde{\mathbf{x}}_0$ to $\tilde{\mathbf{x}}_T$.

   The same construction applies to every curve of C. Moreover, for every $s\in(0,T)$ and every Borel set $B\subseteq\mathbb{R}^D$,

   $$\int_B \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},s)_t\,d\mathbf{x} = p_s^{\mathbf{x}_0,\mathbf{x}_T}(B) = \eta\big(e_s^{-1}(B)\big) = \eta\big(e_s^{-1}(B)\cap C\big),$$

the last equality because $\eta(C) = 1$. The connecting field lines therefore carry the entire flux through every slice $\{t = s\}$, and the remaining ones carry none, which is the zero-flux qualification in (8). These slices are the transversal hypersurfaces used in Appendix B.6.

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

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

where the last equality is the continuity equation for the Conditional Flow. This identity holds pointwise for $t \in (0,T)$.

Fixed total flux and $\Phi_0 = 1$. Let $\varphi \in C_c^\infty(\mathbb{R}^D \times (-\infty, T))$. Integrating the continuity equation by parts on (ε, T) and passing to the limit using Assumption 2 and the weak endpoint condition gives

$$-\int_0^T \int_{\mathbb{R}^D} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot \nabla_{(\mathbf{x},t)}\varphi \, d\mathbf{x}\, dt = \lim_{\varepsilon \to 0^+} \int_{\mathbb{R}^D} p_\varepsilon^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\varphi(\mathbf{x},\varepsilon)\, d\mathbf{x} = \varphi(\mathbf{x}_0, 0).$$

Thus the zero extension satisfies $\nabla \cdot \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} = \delta_{\tilde{\mathbf{x}}_0}$ in $\{t < T\}$ distributionally. For any bounded, piecewise smooth domain M with $\tilde{\mathbf{x}}_0 \in M$ and $\overline{M} \subset \{t < T\}$, the weak flux is therefore

$$\int_{\partial M} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} = (\nabla \cdot \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T})(M) = 1.$$

In the two-sided case, the same calculation at $t = T$ gives $\nabla \cdot \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} = \delta_{\tilde{\mathbf{x}}_0} - \delta_{\tilde{\mathbf{x}}_T}$ on $\mathbb{R}^{D+1}$; the restriction $\overline{M} \subset \{t < T\}$ is then unnecessary if $\tilde{\mathbf{x}}_T \notin \overline{M}$. Hence $\Phi_0 = 1$, independently 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 (16) 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 (Multiplicity of characteristics). The connecting property (8) does not single out one curve, and in general the connecting family is infinite. Consider the two-sided stochastic interpolant of §2.1, with $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}) = \dot{\mathcal{I}}(t,\mathbf{x}_0,\mathbf{x}_T) + (\dot\sigma(t)/\sigma(t))\big(\mathbf{x} - \mathcal{I}(t,\mathbf{x}_0,\mathbf{x}_T)\big)$. The substitution $\mathbf{y} = (\mathbf{x} - \mathcal{I}(t,\mathbf{x}_0,\mathbf{x}_T))/\sigma(t)$ turns the characteristic ODE into $\dot{\mathbf{y}} = 0$, so its solutions form the family

$$\mathbf{x}(t) = \mathcal{I}(t,\mathbf{x}_0,\mathbf{x}_T) + \sigma(t)\,\mathbf{y}, \qquad \mathbf{y} \in \mathbb{R}^D.$$

Since $\sigma(0) = \sigma(T) = 0$, *every* member of this family satisfies $\mathbf{x}(0) = \mathbf{x}_0$ and $\mathbf{x}(T) = \mathbf{x}_T$. The whole continuum of characteristics therefore connects the same pair of endpoints, which is exactly what (8) requires. This is the field-theoretic analogue of a point charge emitting a continuum of field lines. In this example the connecting property holds for every curve, while in general the superposition principle yields it for η-almost every curve, as required by (8).

B.4 Assumptions for Proposition 3.2 and 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 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:
$$\lim_{R\to\infty} \int_{\partial B_R \times [t_1,t_2]} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t)\cdot d\mathbf{S} = 0;$$
3. *(Interior completeness.)* For every $s \in (0,T)$ and every $\mathbf{x} \in \mathbb{R}^D$, the maximal existence interval of the solution of $\dot{\mathbf{x}} = \mathbf{v}_r^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})$, $\mathbf{x}(s) = \mathbf{x}$, with $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T} := \mathbf{E}_{\mathbf{x}}^{\mathbf{x}_0,\mathbf{x}_T}/\mathbf{E}_t^{\mathbf{x}_0,\mathbf{x}_T}$, is all of $(0,T)$.

We use the boundary convention of Appendix B.2.

Assumption 1 already makes $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}$ of class C^1 on $\mathbb{R}^D \times (0,T)$, being a quotient of C^1 functions with nonvanishing denominator. Hence $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}$ is locally Lipschitz in $\mathbf{x}$, uniformly in t on compact subsets of $(0,T)\times\mathbb{R}^D$, so characteristics are unique and depend on initial conditions in a C^1 manner, and no separate Lipschitz assumption is required. This uniqueness holds only up to the maximal existence time, which is why completeness is stated separately in Assumption 3. Together with the connectivity condition of §2.2, which is part of the definition of an Interaction Field rather than an additional requirement, these three assumptions are what the proof in Appendix B.6 uses; the endpoint behaviour of $\mathbf{E}_t^{\mathbf{x}_0,\mathbf{x}_T}$ is derived there and is not assumed.

Remark B.2 (Sufficient criterion for Assumption 3). Assumption 3 holds whenever $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}$ grows at most linearly on compact time intervals: if $\|\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x})\| \le C_{a,b}\,(1+\|\mathbf{x}\|)$ on $[a,b]\times\mathbb{R}^D$ for every $[a,b] \subset (0,T)$, then Grönwall's inequality gives $1+\|\mathbf{x}(t)\| \le (1+\|\mathbf{x}(s)\|)\,e^{C_{a,b}|t-s|}$, so no solution blows up inside $(0,T)$. This criterion gives completeness only, and not the endpoint limits. For the two-sided stochastic interpolant of §2.1 both are available in closed form: the solution through $(s,\mathbf{x})$ is

$$\gamma^{s,\mathbf{x}}(t) = \mathcal{I}(t,\mathbf{x}_0,\mathbf{x}_T) + \frac{\sigma(t)}{\sigma(s)}\big(\mathbf{x} - \mathcal{I}(s,\mathbf{x}_0,\mathbf{x}_T)\big),$$

which is defined on all of $(0,T)$ and, since $\sigma(0) = \sigma(T) = 0$, converges to $\mathcal{I}(0,\mathbf{x}_0,\mathbf{x}_T) = \mathbf{x}_0$ as $t \to 0^+$ and to $\mathcal{I}(T,\mathbf{x}_0,\mathbf{x}_T) = \mathbf{x}_T$ as $t \to T^-$.

B.5 Proof of Proposition 3.2

Here we use Assumptions 1 and 2 above. Fix any $t^* \in (0,T)$. We show that the slice flux equals the total flux:

$$\int_{\mathbb{R}^D} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x},t^*)_t\,d\mathbf{x} = \Phi_0. \tag{28}$$

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

$$\int_{\partial M_R} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} = -\int_{B_R} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, -t^*)_t \, d\mathbf{x} + \int_{\partial B_R \times [-t^*,t^*]} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} + \int_{B_R} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t^*)_t \, d\mathbf{x}. \tag{29}$$

The bottom term vanishes for every R because, by the forward-only convention recorded in Assumption 1, $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t)_t = 0$ for $t < 0$:

$$\int_{B_R} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, -t^*)_t \, d\mathbf{x} = 0.$$

By Assumption 2, the field has vanishing lateral flux at infinity [12]:

$$\lim_{R\to\infty} \int_{\partial B_R \times [-t^*,t^*]} \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T} \cdot d\mathbf{S} = 0,$$

and $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\cdot, t^*)_t$ is integrable, so the top term converges to the integral over $\mathbb{R}^D$. Letting $R \to \infty$ in (29) 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 (28). □

B.6 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. The endpoint behaviour of $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}{}_t$ is not assumed: we derive it from the connectivity condition (8) and Assumptions 1 to 3. Throughout, write

   $$\mu_t(d\mathbf{x}) := \mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t)_t \, d\mathbf{x}, \qquad t \in (0, T),$$

for the slice measure, so that $p_t^{\mathbf{x}_0,\mathbf{x}_T} = \mu_t/\Phi_0$.

(i) Slice-wise form of the connectivity condition. Call a hypersurface $S \subset \mathbb{R}^{D+1}$ transversal if $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}\cdot\mathbf{n} > 0$ along it, and equip it with the flux measure $\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}\cdot\mathbf{n}\, d\mathcal{H}^D$. Since the field is forward-only, ${\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}}_t > 0$, so every field line meets a slice $\{t = s\}$ at most once, and a line connecting $\tilde{\mathbf{x}}_0$ and $\tilde{\mathbf{x}}_T$ meets every intermediate slice exactly once. The slice is thus transversal, and its flux measure is μ_s. Let $A_s \subset \mathbb{R}^D$ be the set of points whose maximal field line, unique by Assumption 1, does not connect $\tilde{\mathbf{x}}_0$ and $\tilde{\mathbf{x}}_T$; it is a Borel set, by the continuity of the flow established in (ii). The flux that these lines carry through $\{t = s\}$ is exactly $\mu_s(A_s)$, so the zero-flux qualification in (8) says that $\mu_s(A_s) = 0$ for every $s \in (0, T)$. Consequently, for μ_s-almost every $\mathbf{x}$, the field line through $(\mathbf{x}, s)$ connects the prescribed endpoints.

(ii) The flow. For $s, t \in (0, T)$ and $\mathbf{x} \in \mathbb{R}^D$ set $X_{s\to t}(\mathbf{x}) := \gamma^{s,\mathbf{x}}(t)$, where $\gamma^{s,\mathbf{x}}$ is the maximal solution of $\dot\gamma(r) = \mathbf{v}_r^{\mathbf{x}_0,\mathbf{x}_T}(\gamma(r))$, $\gamma(s) = \mathbf{x}$. By Assumption 1 the field is C^1 with ${\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}}_t > 0$, so $\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}$ is C^1 and in particular locally Lipschitz in $\mathbf{x}$; characteristics are therefore unique, and Assumption 3 places t in the domain of $\gamma^{s,\mathbf{x}}$. Uniqueness gives $X_{t\to u}\circ X_{s\to t} = X_{s\to u}$ and $X_{s\to s} = \mathrm{id}$, so $X_{s\to t}$ is a bijection of $\mathbb{R}^D$ with inverse $X_{t\to s}$, and C^1 dependence on initial conditions makes it a C^1 diffeomorphism, jointly C^1 in $(t, \mathbf{x})$.

(iii) Transport identity. The divergence-free property (7) reads $\partial_t\mu_t + \nabla_{\mathbf{x}}\cdot(\mathbf{v}_t^{\mathbf{x}_0,\mathbf{x}_T}\mu_t) = 0$ on $(0, T)\times\mathbb{R}^D$. Fix $[a, b] \subset (0, T)$, $s \in [a, b]$ and $\varphi \in C_c^1(\mathbb{R}^D)$, and put $\psi(r, \mathbf{x}) := \varphi\big(X_{r\to s}(\mathbf{x})\big)$, which is defined on all of $\mathbb{R}^D$ by Assumption 3 and is C^1 by (ii); for $r \in [a, b]$ its supports lie in the single compact set $K := \{X_{s\to r}(\mathbf{y}) : r \in [a, b],\ \mathbf{y} \in \operatorname{supp}\varphi\}$. Along any solution γ, uniqueness gives $X_{r\to s}(\gamma(r)) = \gamma(s)$, so $r \mapsto \psi(r, \gamma(r))$ is constant; since a trajectory passes through every point, $\partial_r\psi + \mathbf{v}_r^{\mathbf{x}_0,\mathbf{x}_T}\cdot\nabla_{\mathbf{x}}\psi = 0$. Testing the continuity equation with $\chi(r)\,\psi(r, \mathbf{x})$, $\chi \in C_c^1((a, b))$, therefore leaves only the χ' term:

$$0 = \int_a^b \chi'(r)\, F(r)\, dr, \qquad F(r) := \int \psi(r, \mathbf{x})\, d\mu_r(\mathbf{x}) = \int_K \psi(r, \mathbf{x})\, {\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}}(\mathbf{x}, r)_t\, d\mathbf{x}.$$

The integrand is jointly continuous with support in the fixed compact set K, so F is continuous by dominated convergence, and, having vanishing distributional derivative, it is constant on (a, b) and, by continuity, on $[a, b]$. Evaluating at $r = t$ and $r = s$ gives $(X_{t\to s})_{\#}\mu_t = \mu_s$, that is

$$\mu_t = (X_{s\to t})_{\#}\mu_s, \qquad s, t \in (0, T),$$

since $[a, b] \subset (0, T)$ was arbitrary.

(iv) No exchange of mass at infinity. Every part of μ_t is the image under $X_{s\to t}$ of a part of μ_s, so no mass arrives from infinity, and a push-forward preserves total mass, so none escapes. This excludes simultaneous inflow and outflow at infinity as well as each of them separately.

(v) Endpoint limits. By Proposition 3.2, which rests on Assumptions 1 and 2, each μ_t with $t \in (0, T)$ is a finite measure of mass Φ_0. By (i), for μ_s-a.e. $\mathbf{x}$ the field line through $(\mathbf{x}, s)$ connects $\tilde{\mathbf{x}}_0$ and $\tilde{\mathbf{x}}_T$, and by uniqueness of characteristics that line, reparameterized by t, is the trajectory $r \mapsto X_{s\to r}(\mathbf{x})$ of (ii). Hence $X_{s\to t}(\mathbf{x}) \to \mathbf{x}_0$ as $t \to 0^+$ for μ_s-a.e. $\mathbf{x}$, and for every $\varphi \in C_b(\mathbb{R}^D)$,

$$\int \varphi\, d\mu_t = \int \varphi\big(X_{s\to t}(\mathbf{x})\big)\, d\mu_s(\mathbf{x}) \ \longrightarrow\ \Phi_0\, \varphi(\mathbf{x}_0), \qquad t \to 0^+,$$

by dominated convergence, using $|\varphi| \le \|\varphi\|_\infty$ and finiteness of μ_s. The terminal condition follows symmetrically. In the sense of weak convergence of measures on $\mathbb{R}^D$, therefore,

$${\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}}(\mathbf{x}, t)_t\, d\mathbf{x} \ \Rightarrow\ \Phi_0\, \delta_{\mathbf{x}_0}(\mathbf{x})\ (t \to 0^+), \quad {\mathbf{E}^{\mathbf{x}_0,\mathbf{x}_T}}(\mathbf{x}, t)_t\, d\mathbf{x} \ \Rightarrow\ \Phi_0\, \delta_{\mathbf{x}_T}(\mathbf{x})\ (t \to T^-).$$

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),

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

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 (22).

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 Additional material for Section 4

This section collects the assumptions and proofs for the results of §4.

C.1 Assumptions for Proposition 4.1

Write $\mathbf{b} := -\mathbf{E}$ for the generative field and assume that its trajectories issued from $U \times \{T\}$ satisfy:

1. They are unique and sweep out an open subset of $\mathbb{R}^{D+1}$, on which $\mathbf{b}$ is C^1 and nonzero away from the initial and terminal sections, with C^1 one-sided traces at charged sections;
2. They leave $U \times \{T\}$ into $t < T$ transversally: $b_t(a, T^-) \neq 0$ for $a \in U$;
3. They reach the selected endpoint on $t = 0$ from $t < 0$ transversally: $b_t(F(a), 0^-) \neq 0$ for $a \in U$;
4. Any crossing of a plate before the selected endpoint lies outside the charged support, so that $\mathbf{b}$ extends C^1 across it.

We call such a family a C^1 *flow tube*. The implicit-function theorem makes its hitting time $\tau(a)$ and transport map F of class C^1 on U, and both endpoint sections are oriented using the same data coordinates $(x_1, \dots, x_D)$. These assumptions constrain only the branch under consideration and not the field away from it; Lemma C.2 verifies them for the EFM construction of Proposition 4.2.

C.2 Proof of Proposition 4.1

Let φ_τ denote the flow of the generative field $\mathbf{b} = -\mathbf{E}$, and let $e_1, \dots, e_D$ be the positively oriented coordinate basis tangent to the plates. Fix $a \in U$ and let $\tau(a)$ be the hitting time of the selected intersection, so that

$$\varphi_{\tau(a)}(a, T) = (F(a), 0).$$

We show that the orientation of F is decided by the normal traces of $\mathbf{b}$ at the two ends of the tube. The argument rests on the standard Jacobian identity for a flow map between two transverse sections, which we derive here for completeness:

$$b_t(F(a), 0^-) \det \nabla F(a) = b_t(a, T^-) \det D\varphi_{\tau(a)}. \tag{30}$$

First, let's build a frame that joins the two sections. Let $M(a)$ be the $(D+1) \times (D+1)$ matrix whose first D columns are the variations $\partial_{a_i}\big[\varphi_{\tau(a)}(a, T)\big]$ and whose last column is $\mathbf{b}(F(a), 0^-)$. We evaluate $\det M(a)$ in two ways.

At the terminal section. Since $\varphi_{\tau(a)}(a, T) = (F(a), 0)$ for every $a \in U$, the first D columns are $(\partial_i F(a), 0)$, so expansion along the last row gives

$$\det M(a) = b_t(F(a), 0^-) \det \nabla F(a).$$

Along the flow. Differentiating $\varphi_{\tau(a)}(a, T)$ with respect to a produces the fixed-time variations $D\varphi_{\tau(a)} e_i$ and an additional multiple of $\mathbf{b}(F(a), 0^-)$ coming from the derivative of $\tau(a)$; being multiples of the last column, they leave the determinant unchanged. Moreover, invariance of a vector field under its own flow gives

$$D\varphi_{\tau(a)}\, \mathbf{b}(a, T^-) = \mathbf{b}(F(a), 0^-),$$

so the last column is the image of $\mathbf{b}(a, T^-)$ as well. Hence

$$\det M(a) = \det D\varphi_{\tau(a)} \cdot \det\big(e_1, \dots, e_D, \mathbf{b}(a, T^-)\big) = b_t(a, T^-) \det D\varphi_{\tau(a)}.$$

Comparing the two expressions proves (30).

Next, let's check that the flow factor is positive. By item **1**, the field $\mathbf{b}$ is C^1 along the tube, including the one-sided traces at its two ends and the crossing in between, so φ_τ is a diffeomorphism onto its image for every $\tau \in [0, \tau(a)]$. Therefore $\tau \mapsto \det D\varphi_\tau$ is continuous, never vanishes, and equals 1 at $\tau = 0$, whence $\det D\varphi_{\tau(a)} > 0$.

It remains to read off the two traces. Since the tube leaves $U \times \{T\}$ into $t < T$, we have $b_t(a, T^-) < 0$; since it reaches the selected endpoint from $t < 0$, we have $b_t(F(a), 0^-) > 0$. The traces have opposite signs, so (30) gives $\det \nabla F(a) < 0$ for every $a \in U$.

The same computation applies to a branch reaching the target plate from $t > 0$: it replaces $b_t(F(a), 0^-)$ by $b_t(F(a), 0^+) < 0$ and gives $\det \nabla F(a) > 0$. The reversal is thus caused by the return from $t < 0$, and not by nonmonotone motion inside the slab. □

C.3 Assumptions for Proposition 4.2

We use the standard EFM convention in $\mathbb{R}^{D+1}$:

1. For a positive constant c_D, the Coulomb kernel is $K(y) := c_D\, y/\|y\|_2^{D+1}$, and the global field of target and source densities π_0 and π_T, carried by the plates $t = 0$ and $t = T$ respectively, is

$$\mathbf{E}_{\pi_0,\pi_T}(x,t) := \int_{\mathbb{R}^D} K\big((x,t)-(y,0)\big)\pi_0(y)\,dy - \int_{\mathbb{R}^D} K\big((x,t)-(y,T)\big)\pi_T(y)\,dy; \tag{31}$$

2. The first plate carries positive charge and the second carries negative charge, so generative trajectories follow $\mathbf{b}_{\pi_0,\pi_T} := -\mathbf{E}_{\pi_0,\pi_T}$. This is the setting of Kolesov et al. [11], Manukhov et al. [12] in the notation of §2.2, with z renamed t and L renamed T: their field is our $\mathbf{E}$, their direction of motion is our $\mathbf{b}$, and the one-sided normal traces satisfy

$$E_z(x,s^{\pm}) = \mathbf{E}_t(x,s^{\pm}) = -b_t(x,s^{\pm}), \qquad s \in \{0,T\};$$

3. Trajectories follow the stochastic map of Manukhov et al. [12, Appendix A.2, Eqs. (25)–(27)] in these coordinates, including its convention of initializing a trajectory arbitrarily close to a plate.

C.4 Proof of Proposition 4.2

We first record the one-sided traces of a Coulomb layer and then construct densities whose field carries a returning tube.

Lemma C.1 (One-sided traces of a Coulomb layer)**.** *Let $\sigma \in C_c^\infty(\mathbb{R}^D)$ be a positive surface density on the plate $t = s$, and let*

$$\mathbf{E}_\sigma(x,t) = c_D \int_{\mathbb{R}^D} \frac{(x-y,t-s)}{\big(\|x-y\|_2^2+(t-s)^2\big)^{(D+1)/2}}\sigma(y)\,dy.$$

Then $\mathbf{E}_\sigma$ is smooth off the plate and has smooth one-sided traces. Its normal component satisfies

$$(E_\sigma)_t(x,s^+) = \kappa_D\sigma(x), \qquad (E_\sigma)_t(x,s^-) = -\kappa_D\sigma(x), \tag{32}$$

where

$$\kappa_D := c_D \int_{\mathbb{R}^D} \frac{du}{(1+\|u\|_2^2)^{(D+1)/2}} > 0.$$

The tangential component has the same principal-value trace from both sides. Consequently, if a field H is smooth across the plate, then adding the layer σ gives normal traces $H_t(x,s) \pm \kappa_D\sigma(x)$; for a negative layer $-\sigma$, the signs are reversed.

Proof. Writing $h = t - s$ and changing variables $y = x + |h|u$ in the normal component gives

$$(E_\sigma)_t(x,s+h) = c_D\,\mathrm{sgn}(h) \int_{\mathbb{R}^D} \frac{\sigma(x+|h|u)}{(1+\|u\|_2^2)^{(D+1)/2}}\,du.$$

Dominated convergence proves (32). The tangential statement is the standard principal-value limit of the odd Coulomb kernel; for a smooth compactly supported density this singular integral is smooth. □

Lemma C.2 (Persistence of a returning dipole tube)**.** *For every $D \geq 1$, there exist a radial density $\rho \in C_c^\infty(\mathbb{R}^D)$, a point $a_\star \in \mathbb{R}^D \setminus \{0\}$, and, for every sufficiently small $\varepsilon > 0$, smooth probability densities*

$$\pi_0^\varepsilon(x) = \rho_\varepsilon(x), \qquad \pi_T^\varepsilon(x) = (1-\delta_\varepsilon)\rho_\varepsilon(x) + \delta_\varepsilon\rho_\varepsilon(x-a_\star),$$

where $\rho_\varepsilon(x) = \varepsilon^{-D}\rho(x/\varepsilon)$ and $\delta_\varepsilon = \varepsilon^{D+2}$, with the following property. The corresponding EFM field has a nonempty open set U_ε satisfying

$$\overline{U_\varepsilon} \subset a_\star + \varepsilon\,\mathrm{int}(\mathrm{supp}\,\rho)$$

such that every generative trajectory issued from U_ε leaves the source plate into the slab, crosses $t = 0$ outside $\operatorname{supp}\pi_0^\varepsilon$, enters $t < 0$, and returns to $t = 0$ inside $\operatorname{int}(\operatorname{supp}\pi_0^\varepsilon)$. These trajectories form a C^1 flow tube in the sense of Appendix C.1. On U_ε and its terminal footprint V_ε, respectively, the one-sided traces satisfy

$$b_t(a, T^-) < 0, \qquad b_t(a, T^+) < 0, \qquad b_t(x, 0^-) > 0 > b_t(x, 0^+).$$

Proof. We construct a transverse returning tube for two point charges and then smooth the charges, treating in turn the middle of the tube, its departure from the source plate, its passage through the annulus around the target support, and its capture by the target plate.

A reference dipole. Put a positive unit charge at $q_+ = (0, 0)$ and a negative unit charge at $q_- = (0, T)$. Away from these two points, their field is

$$\mathbf{E}_\star(x, t) = c_D \left[\frac{(x, t)}{(\|x\|_2^2 + t^2)^{(D+1)/2}} - \frac{(x, t - T)}{(\|x\|_2^2 + (t - T)^2)^{(D+1)/2}} \right]. \tag{33}$$

This field is rotationally symmetric in x, and its field lines lie in two-dimensional meridional half-planes. It has no zero away from the two charges: equality of the two terms in (33) would imply

$$\frac{y - q_+}{\|y - q_+\|_2^{D+1}} = \frac{y - q_-}{\|y - q_-\|_2^{D+1}},$$

which is impossible because the radial map $z \mapsto z/\|z\|_2^{D+1}$ is injective. Its potential $\psi_\star$, defined by $\mathbf{E}_\star = -\nabla\psi_\star$, is the superposition of $\pm g(\|y - q_\pm\|_2)$ with

$$g(r) := \begin{cases} -c_1 \log r, & D = 1, \\ \dfrac{c_D}{D - 1}\, r^{-(D-1)}, & D \geq 2, \end{cases} \tag{34}$$

so that g is strictly decreasing in every dimension.

The positive charge emits positive flux into the lower half-space. Since the two charges have zero net charge, $\|\mathbf{E}_\star(y)\| = O(\|y\|^{-D-1})$ at infinity, so its absolute flux through the sphere of radius R is $O(R^{-1})$. Call a field line *escaping* if it meets spheres of arbitrarily large radius about the origin; a line that is not escaping stays in a bounded set. The flux carried by a tube of field lines is constant because $\nabla \cdot \mathbf{E}_\star = 0$ away from the charges, so a tube of fixed positive flux cannot consist of escaping lines. Every angular neighborhood of the lower axis carries a fixed positive flux, and therefore contains a line that stays in a bounded set.

Let γ be such a line, emitted into $t < 0$ and traced in the direction of the field. Since $\mathbf{E}_\star = -\nabla\psi_\star$, the potential is strictly decreasing along γ, at rate $-\|\mathbf{E}_\star\|_2^2 < 0$. Its ω-limit set is nonempty and invariant, and $\psi_\star$ is constant on it; hence it consists of zeros of $\mathbf{E}_\star$ or of the two charges. There are no zeros away from $q_\pm$, and q_+ is excluded because $\psi_\star \to +\infty$ there whereas $\psi_\star$ decreases along γ. Thus γ terminates at q_- and, in particular, crosses $t = 0$ on its way. Returning field lines therefore accumulate at the lower axial line.

Reflection in the midplane, $R(x, t) = (x, T - t)$, satisfies $\mathbf{E}_\star(Ry) = -DR\,\mathbf{E}_\star(y)$ and hence maps unoriented field lines to unoriented field lines. A connecting line crosses the fixed plane $t = T/2$; uniqueness of integral curves makes that line invariant under R. Its off-axis crossings of the two plates are therefore paired. For $x \neq 0$,

$$(E_\star)_t(x, 0) = (E_\star)_t(x, T) = c_D \frac{T}{(\|x\|_2^2 + T^2)^{(D+1)/2}} > 0,$$

so both crossings are transverse. Since the normal component has only the upward sign at either plate, the segment between these crossings remains in the open slab. We choose below a non-axial returning line close to the lower axis and denote its consecutive crossings by $(x_\star, 0)$ and $(a_\star, T)$; being non-axial, it has $x_\star \neq 0$ and $a_\star \neq 0$. Traced in the generative direction $\mathbf{b}_\star = -\mathbf{E}_\star$, it starts at $(a_\star, T)$, stays in $0 < t < T$ until $(x_\star, 0)$, enters $t < 0$, and returns to q_+. A sufficiently small tube around it has the same transverse crossings.

Smooth charges and the middle of the tube. Choose a radially symmetric function $\rho \in C_c^\infty(\mathbb{R}^D)$, supported in the unit ball, positive in its interior, strictly decreasing in radius for $0 < \|x\|_2 < 1$, and

normalized by $\int \rho = 1$. Fix $0 < r < 1$ with $\overline{B_r} \subset \{\rho > 0\}$. We choose a large fixed L, a small fixed $r_0 > 0$, and then the reference tube as specified by the annular estimates below. Set

$$\rho_\varepsilon(x) := \varepsilon^{-D}\rho(x/\varepsilon)$$

and couple the two small parameters by

$$\delta_\varepsilon := \varepsilon^{D+2}.$$

For sufficiently small $\varepsilon > 0$, define

$$\pi_0^\varepsilon(x) := \rho_\varepsilon(x), \tag{35}$$
$$\pi_T^\varepsilon(x) := (1-\delta_\varepsilon)\rho_\varepsilon(x) + \delta_\varepsilon\rho_\varepsilon(x - a_\star). \tag{36}$$

Both densities are smooth, compactly supported probability densities. For small ε, the two source components have disjoint supports.

Fix a compact sub-tube K around the middle part of the reference tube. We truncate it inside a fixed neighborhood of $(a_\star, T)$ and at the sphere $\|(x,t)\|_2 = r_0$ around q_+, while retaining its first crossing $(x_\star, 0)$. Standard convolution estimates away from the three shrinking supports give

$$\left\|\mathbf{E}_{\pi_0^\varepsilon,\pi_T^\varepsilon} - \mathbf{E}_\star\right\|_{C^1(K)} \longrightarrow 0. \tag{37}$$

Continuous dependence of ODE solutions therefore preserves the open tube, the order of its crossings, and their transversality on K. In particular, its first crossing of $t = 0$ remains near $x_\star$ and outside $\operatorname{supp}\pi_0^\varepsilon$. The field is smooth across the plate there, so the EFM stopping probability at this crossing is zero.

Extension to the source plate. After taking one-sided traces, the field and first derivative generated by the satellite bump obey

$$\|E_{\mathrm{sat}}\| = O(\delta_\varepsilon\varepsilon^{-D}) = O(\varepsilon^2), \qquad \|DE_{\mathrm{sat}}\| = O(\delta_\varepsilon\varepsilon^{-D-1}) = O(\varepsilon).$$

Thus they are a C^1-small one-sided perturbation of the reference field in a fixed neighborhood of $(a_\star, T)$. Since $(E_\star)_t(a_\star, T) > 0$, Lemma C.1 and continuity give a nonempty open set

$$\overline{U_\varepsilon} \subset a_\star + \varepsilon\,\mathrm{int}(\operatorname{supp}\rho)$$

and $c_s > 0$ such that

$$\mathbf{E}_t(a, T^\pm) \geq c_s, \qquad b_t(a, T^\pm) \leq -c_s, \qquad a \in U_\varepsilon.$$

The generative branch from each $a \in U_\varepsilon$ therefore leaves into the slab and joins the perturbed middle tube.

Passage through the annulus. We bridge the fixed middle tube and the $O(\varepsilon)$ neighborhood of the target support. Write $y = (x,t) = \varrho\omega$, where $\varrho = \|y\|_2$ and $\omega \in S^D$, and decompose the generative field in $B_{r_0}(q_+)$ as

$$\mathbf{b}(y) = -c_D\varrho^{-D}\omega + \mathbf{p}(y).$$

Taylor expansion of the Coulomb kernel under convolution with ρ_ε, whose first moment vanishes by radial symmetry, together with the $O(1)$ bound on the field of the source plate near q_+, gives, for $L\varepsilon \leq \varrho \leq r_0$,

$$e(\varrho) := \sup_{\|y\|_2=\varrho} \frac{\varrho^D\|\mathbf{p}(y)\|_2}{c_D} \leq C_1\left(\frac{\varepsilon}{\varrho}\right)^2 + C_2\varrho^D, \tag{38}$$

where C_1, C_2 are independent of ε, $L \geq 2$, and $r_0 \leq T/2$ (and may be chosen uniformly for reference lines in a fixed small cone about the lower axis). If $e \leq \frac{1}{2}$, then $\mathbf{b}\cdot\omega \leq -\frac{1}{2}c_D\varrho^{-D} < 0$, so ϱ decreases strictly along trajectories, and

$$\left\|\frac{d\omega}{d\varrho}\right\|_2 = \frac{\|(I - \omega\omega^\top)\mathbf{b}\|_2}{\varrho|\mathbf{b}\cdot\omega|} \leq \frac{2e(\varrho)}{\varrho}.$$

Thus, between radii r_0 and $L\varepsilon$, the direction ω changes by at most

$$2\int_{L\varepsilon}^{r_0} e(\varrho)\,\frac{d\varrho}{\varrho} \leq C_1L^{-2} + \frac{2C_2}{D}r_0^D.$$

To make the choices above explicit, choose L large enough that $C_1 L^{-2} < \frac{1}{4}\arctan(r/L)$ and set $\Theta := \frac{1}{4}\arctan(r/L)$. Then choose r_0 so that $2C_2 r_0^D/D < \Theta$ and $e \le \frac{1}{2}$ on the annulus, and finally the reference tube so that its angle with the lower axis at radius r_0 is less than Θ; such a line exists because the returning lines of $\mathbf{E}_\star$ accumulate at the lower axial line. For sufficiently small ε, the perturbed tube enters $B_{r_0}(q_+)$ within angle 2Θ of the lower axis and therefore stays within angle $4\Theta = \arctan(r/L)$ of it throughout the annulus. In particular, it stays in $t < 0$ and crosses the section

$$\{(x,t) : t = -L\varepsilon,\ \|x\|_2 < r\varepsilon\}. \tag{39}$$

Capture by the target layer. Put $(x,t) = \varepsilon(\xi, s)$ and multiply the generative field by ε^D. On compact subsets of $s \le 0$, using the one-sided trace at $s = 0$,

$$\varepsilon^D \mathbf{b}_{\pi_0^\varepsilon, \pi_T^\varepsilon}(\varepsilon\xi, \varepsilon s) \longrightarrow \mathbf{b}_\rho(\xi, s) \quad \text{in } C^1, \tag{40}$$

where $\mathbf{b}_\rho$ is the reversed field of the single positive layer ρ at $s = 0$. The target-layer term rescales exactly, whereas the two source components vanish after multiplication by ε^D because their supports remain a positive distance away. For $s < 0$,

$$(b_\rho)_s(\xi, s) = c_D(-s) \int_{\mathbb{R}^D} \frac{\rho(\eta)\, d\eta}{(\|\xi - \eta\|_2^2 + s^2)^{(D+1)/2}} > 0.$$

The potential of the layer ρ at a fixed height $s \le 0$ is $\psi_\rho(\xi, s) = \int_{\mathbb{R}^D} \rho(\eta) G_s(\xi - \eta)\, d\eta$ with $G_s(u) := g\big((\|u\|_2^2 + s^2)^{1/2}\big)$ and g as in (34); thus G_s is radial and strictly decreasing in $\|u\|_2$. Write $\rho(\eta) = \hat\rho(\|\eta\|_2)$, fix $r > 0$, and take $\xi = re_1$ by rotational symmetry. Integrating by parts in η_1 and pairing η with its reflection $\eta^* := (-\eta_1, \eta_2, \dots, \eta_D)$ gives

$$\partial_r \psi_\rho(re_1, s) = \int_{\{\eta_1 > 0\}} \hat\rho'(\|\eta\|_2)\, \frac{\eta_1}{\|\eta\|_2} \big[G_s(re_1 - \eta) - G_s(re_1 - \eta^*)\big]\, d\eta,$$

which is valid up to $s = 0$, where G_s is only weakly singular, because the derivative falls on ρ. For $\eta_1 > 0$ one has $\|re_1 - \eta\|_2 < \|re_1 - \eta^*\|_2$, so the bracket is positive, while $\hat\rho' \le 0$ and $\hat\rho' < 0$ on a set of positive measure because $\hat\rho$ is strictly decreasing on $(0,1)$. Hence $\partial_r \psi_\rho < 0$ for every $r > 0$ and $s \le 0$, and since $\mathbf{b}_\rho$ is the gradient of ψ_ρ,

$$(b_\rho)_\xi(\xi, s) \cdot \xi = \|\xi\|_2\, \partial_r \psi_\rho(\|\xi\|_2, s) < 0 \qquad (\xi \neq 0,\ s \le 0),$$

where the one-sided trace is used at $s = 0$. Consider the vertical cylinder

$$\mathcal{C}_L := \{(\xi, s) : \|\xi\|_2 < r,\ -L < s < 0\}.$$

On its compact lateral boundary and closure, respectively, there are constants $c_1, c_0 > 0$ such that

$$(b_\rho)_\xi(\xi, s) \cdot \xi \le -c_1, \qquad (b_\rho)_s(\xi, s) \ge c_0.$$

For the second inequality, use positivity for $s < 0$ and the jump trace $(b_\rho)_s(\xi, 0^-) = \kappa_D \rho(\xi) > 0$ on $\overline{B_r}$. Hence the limiting field points strictly inward on the lateral boundary, and every trajectory entering through the bottom remains in the cylinder and reaches $s = 0$ in finite flow time. By (40), the same strict inequalities hold for the rescaled full field when ε is small. Together with (39), this carries the perturbed tube to $t = 0$ inside a compact subset of $\mathrm{int}(\mathrm{supp}\, \pi_0^\varepsilon)$.

Let V_ε be its terminal footprint. The target-layer traces have order ε^{-D} there, while the field of the distant source components is $O(1)$. Lemma C.1 therefore gives, uniformly on V_ε,

$$b_t(x, 0^-) > 0 > b_t(x, 0^+)$$

for sufficiently small ε. The terminal crossing is transverse, and the resulting trajectories form the claimed C^1 flow tube. □

We can now prove the proposition.

First, let's fix the geometry. Choose the densities and $U = U_\varepsilon$ from Lemma C.2; the trajectories issued from U form a C^1 flow tube satisfying the assumptions of Appendix C.1.

Next, let's evaluate the probabilities of the stochastic map [12, Appendix A.2, Eqs. (25)–(27)] in the coordinates of the setup. At the source plate, Lemma C.2 gives $b_t(a, T^\pm) < 0$: both one-sided

generative traces point into the slab, so no branch leaves into $t > T$ and the branch entering $t < T$ is selected with probability $\mu = 1$. At the first target crossing, which lies outside $\operatorname{supp} \pi_0^\varepsilon$, the field is continuous, so the one-sided normal components agree and the stopping probability is $\nu = 0$. At the second crossing, $b_t(x, 0^-) > 0 > b_t(x, 0^+)$, so the one-sided normal components have opposite signs and $\nu = 1$. Hence a trajectory started in U follows the returning tube and terminates at its second intersection with probability one. The main and satellite source supports are disjoint for small ε because $a_\star \neq 0$, and hence

$$\pi_T^\varepsilon(U) = \delta_\varepsilon \int_U \rho_\varepsilon(x - a_\star)\, dx > 0,$$

so the returning event has positive probability.

It remains to orient the resulting transport map F. The traces of Lemma C.2, $b_t(a, T^-) < 0$ and $b_t(F(a), 0^-) > 0$, are exactly the transversality conditions of Appendix C.1, so Proposition 4.1 gives $\det \nabla F(a) < 0$ for every $a \in U$. In particular, the reversal is not an artifact of a singular branch: the field is smooth and nonzero along the tube away from the two charged plates, and F is a local C^1 diffeomorphism on U. □

D Empirical Validations

We first describe the toy-data setup and then examine the three takeaways of §5: flow–field equivalence, the probabilistic interpretation of fields, and field-informed volume coverage.

D.1 Setup for the toy-data validation

We describe the common setup shared by the three verifications of Appendices D.2–D.4, which support the takeaways of §5. All of them are performed on the two-dimensional toy dataset shown in Figure 2.

We use two two-sided stochastic interpolants (§2.1) and set $T = 1$ throughout. For $i \in \{1, 2\}$ and $t \in [0, 1]$, define

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

Equivalently,

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

The first interpolant uses a linear path and the second a trigonometric one,

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

while both use the same noise schedule $\sigma^1(t) = \sigma^2(t) = 0.05\sqrt{t(1-t)}$.

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

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

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 the global field obtained by the generalized superposition principle (9),

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

These two interpolants define the matched CFM–IFM pairs used throughout Appendices D.2–D.4.

Unless stated otherwise, training points are drawn as $t \sim \mathcal{U}[0, T]$ and $\mathbf{x} \sim p_{i,t}^{\mathbf{x}_0, \mathbf{x}_T}$, so IFM is trained with the volume distribution induced by the field itself, $p_{\text{vol}} = p_{i,t}^{\mathbf{x}_0, \mathbf{x}_T}$, and both models compared in Appendix D.2 see points from the same distribution.

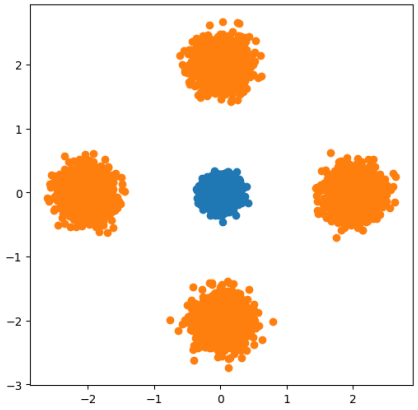

Figure 2: Two-dimensional toy Gaussian dataset used for the empirical validation. Blue points denote samples from the source distribution π_T, a narrow Gaussian centered near the origin. Orange points denote samples from the target distribution π_0, a symmetric four-component Gaussian mixture with modes arranged along the coordinate axes.

D.2 Empirical verification of the duality

Equations (16) and (21) make the equivalence half of Takeaway 1 of §5 directly testable: they map a conditional flow to a forward-only interaction field and back, so the two descriptions of one interpolant must induce the same data-space dynamics.

For each interpolant $i \in \{1, 2\}$ of Appendix D.1 we train two separate models, each with its own objective and its own target construction. The CFM model learns the drift $\mathbf{v}_i$ from the conditional flow $(\mathbf{v}_{i,t}^{\mathbf{x}_0,\mathbf{x}_T}, p_{i,t}^{\mathbf{x}_0,\mathbf{x}_T})$ with the CFM objective (5); the IFM model learns the global field $\mathbf{E}_i$ from the conditional field $\mathbf{E}_i^{\mathbf{x}_0,\mathbf{x}_T}$ with the normalized-field objective (13).

We then compare the data-space velocities the two models induce: the CFM drift $\mathbf{v}_i$, and the drift read off the trained field (22),

$$\mathbf{u}_i(\mathbf{x}, t) := \frac{\mathbf{E}_i(\mathbf{x}, t)_{\mathbf{x}}}{\mathbf{E}_i(\mathbf{x}, t)_t},$$

which is a ratio of components and therefore unaffected by the normalization of the field. If the duality holds, a matched pair $(\mathbf{v}_i, \mathbf{u}_i)$, both sides from the same interpolant, must describe the same dynamics, while a mismatched pair, from different interpolants, must not.

We report the mean squared error between the two drifts and the relative error

$$\mathrm{RelMSE}(\mathbf{v}, \mathbf{u}) = \mathbb{E}\left[\|\mathbf{u}(\mathbf{x}, t) - \mathbf{v}(\mathbf{x}, t)\|_2^2 \,/\, \|\mathbf{v}(\mathbf{x}, t)\|_2^2\right], \tag{41}$$

where the first listed drift of each comparison serves as the reference $\mathbf{v}$. Both drifts are evaluated at the same $(\mathbf{x}, t)$, obtained by deterministic linear interpolations between sampled endpoints.

Table 3 confirms this prediction. The matched MSEs are of order 10^{-7}, six orders of magnitude below the mismatched ones, and the matched relative errors four to five orders below: $\mathbf{v}_1$ agrees with $\mathbf{u}_1$, and $\mathbf{v}_2$ agrees with $\mathbf{u}_2$, whereas dynamics built from different interpolants remain far apart. The two representations therefore coincide when they come from the same interpolant, while different interpolants define genuinely different data-space dynamics. Since $\mathbf{u}_i$ and $\mathbf{v}_i$ are the same drift, comparing the two fields and comparing the two flows is the same measurement, which is why the two mismatched rows coincide.

D.3 Probabilistic interpretation of a forward-only field

We illustrate Takeaway 2 of §5 using the first interpolant. We evolve particles under the dynamics induced by the IFM model trained in Appendix D.2 and record their empirical distributions at several times. We compare these distributions with heatmaps of the last component $\mathbf{E}_1(\mathbf{x}, t)_t$ of the global field, estimated by Monte Carlo averaging the conditional fields over sampled endpoint pairs. As shown in Figure 3, the transported particles concentrate in regions where this component is large, supporting the probabilistic interpretation predicted by the flow–field duality.

Table 3: MSE and RelMSE between the CFM drift and the IFM-induced drift, for pairs built from the same interpolant (matched) and from different interpolants (mismatched). Here $\mathbf{u}_i = \mathbf{E}_i(\mathbf{x}, t)_{\mathbf{x}}/\mathbf{E}_i(\mathbf{x}, t)_t$ and RelMSE is defined in (41).

Type	**Comparison**	**MSE** (↓)	**RelMSE** (↓)
Matched	$\mathbf{v}_1$ vs. $\mathbf{u}_1$	$\mathbf{2.05 \times 10^{-7}}$	$\mathbf{1.08 \times 10^{-6}}$
Matched	$\mathbf{v}_2$ vs. $\mathbf{u}_2$	$\mathbf{2.84 \times 10^{-7}}$	$\mathbf{1.34 \times 10^{-5}}$
Mismatched	$\mathbf{v}_1$ vs. $\mathbf{v}_2$	4.93×10^{-1}	0.34
Mismatched	$\mathbf{u}_1$ vs. $\mathbf{u}_2$	4.93×10^{-1}	0.34

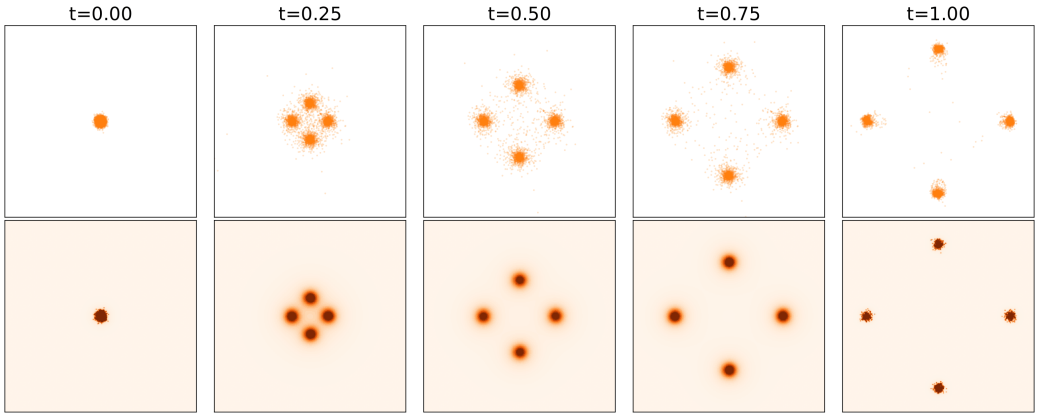

Figure 3: Probabilistic interpretation of a forward-only interaction field. Top: particles evolved under the dynamics induced by the trained IFM model. Bottom: heatmaps of the last component $\mathbf{E}_1(\mathbf{x}, t)_t$ of the global field, estimated by Monte Carlo averaging the conditional fields over sampled endpoint pairs.

D.4 Field-informed volume coverage

To assess whether the field-induced density of Takeaway 3 of §5 provides a useful training distribution, we perform a matched–mismatched ablation. For each field $\mathbf{E}_i$, we train the normalized-field objective (13) with volume distribution p_{vol} set either to the matched path $p_{i,t}^{\mathbf{x}_0,\mathbf{x}_T}$, the conditional density read off the same field as $\mathbf{E}_i^{\mathbf{x}_0,\mathbf{x}_T}(\mathbf{x}, t)_t/\Phi_0$, which reproduces (24), or to the mismatched path $p_{j,t}^{\mathbf{x}_0,\mathbf{x}_T}$ read off the other field, $j \neq i$. Thus, the objective and the target field are fixed, and only the volume-coverage distribution is changed.

Table 4: Ablation of the volume-coverage distribution. For each field $\mathbf{E}_i$, the normalized-field objective (13) is trained with p_{vol} set to the matched path $p_{i,t}^{\mathbf{x}_0,\mathbf{x}_T}$ or to the mismatched path $p_{j,t}^{\mathbf{x}_0,\mathbf{x}_T}$, $j \neq i$; the last column reports the final W_2 distance. Reported as mean ± standard deviation over 3 repeated evaluations of each trained model, with randomly resampled source subsets.

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 4 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.

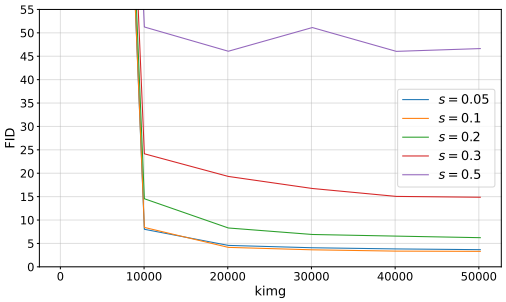

Figure 4: FID vs. training progress for the two-sided interpolant at different noise scales s.

E Additional Discussion and Experiments on Duality

Takeaways 2 and 3 of §5 are mostly about what CFM lends IFM: an explicit marginal path, and a volume distribution that no longer has to be guessed. The duality also runs the other way, and what IFM lends CFM is its estimator. By Theorem 3.1 the global drift of a conditional flow is the ratio of the components of its dual field, and that field is estimated by superposition over endpoint pairs (9). Read back in data space, the single-sample CFM target (5) becomes a multi-sample one,

$$\hat{\mathbf{v}}_t^N(\mathbf{x}) := \frac{\hat{\mathbf{E}}^N(\mathbf{x},t)_{\mathbf{x}}}{\hat{\mathbf{E}}^N(\mathbf{x},t)_t} = \frac{\sum_{i=1}^N \mathbf{v}_t^{\mathbf{x}_0^i,\mathbf{x}_T^i}(\mathbf{x})\, p_t^{\mathbf{x}_0^i,\mathbf{x}_T^i}(\mathbf{x})}{\sum_{i=1}^N p_t^{\mathbf{x}_0^i,\mathbf{x}_T^i}(\mathbf{x})} = \sum_{i=1}^N \omega_i\, \mathbf{v}_t^{\mathbf{x}_0^i,\mathbf{x}_T^i}(\mathbf{x}), \tag{42}$$

with $\omega_i := p_t^{\mathbf{x}_0^i,\mathbf{x}_T^i}(\mathbf{x}) / \sum_{k=1}^N p_t^{\mathbf{x}_0^k,\mathbf{x}_T^k}(\mathbf{x})$ and $\mathbf{x}_0^i, \mathbf{x}_T^i \sim \pi_{0,T}$. This appendix asks whether the borrowed estimator helps conditional flows in practice.

ExFM [21] derives density-weighted targets for noisy two-sided Gaussian paths with arbitrary endpoint couplings; related estimators appear in [22, 23]. Here, the same type of target follows from superposition of the dual interaction fields.

E.1 Setup

Two-sided interpolant. Following [16], it 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 carries an extra $\sqrt{2}$, we report results in terms of $s := C/\sqrt{2}$. Figure 4 shows that larger s consistently degrades sample quality, so we use $s = 0.1$ throughout. The network regresses the conditional velocity of Table 2 through (42).

One-sided interpolants. For EDM [19] and PFGM++ [10] the conditional velocities are also given in Table 2. Following PFGM++ [10], we parameterize the auxiliary radius as $r = \sqrt{d}\,\sigma$. The weights in (42) are computed from $p_\sigma(\mathbf{x}_\sigma \mid \mathbf{x}_0^i)$ for $\mathbf{x}_0^i \sim \pi_0$. To keep the original hyperparameters and samplers, we convert the velocity estimate into a clean-sample one:

$$\begin{aligned}
\mathbf{x}_0^N(\mathbf{x}_\sigma) &= \mathbf{x}_\sigma - \sigma\, \hat{\mathbf{v}}_\sigma^N(\mathbf{x}_\sigma) \\
&= \mathbf{x}_\sigma - \sigma\, \frac{\sum_{i=1}^N \mathbf{v}_\sigma^{\mathbf{x}_0^i}(\mathbf{x}_\sigma)\, p_\sigma(\mathbf{x}_\sigma \mid \mathbf{x}_0^i)}{\sum_{i=1}^N p_\sigma(\mathbf{x}_\sigma \mid \mathbf{x}_0^i)} \\
&= \mathbf{x}_\sigma - \frac{\sum_{i=1}^N (\mathbf{x}_\sigma - \mathbf{x}_0^i)\, p_\sigma(\mathbf{x}_\sigma \mid \mathbf{x}_0^i)}{\sum_{i=1}^N p_\sigma(\mathbf{x}_\sigma \mid \mathbf{x}_0^i)},
\end{aligned} \tag{43}$$

where the last equality uses $\mathbf{v}_\sigma^{\mathbf{x}_0^i}(\mathbf{x}_\sigma) = (\mathbf{x}_\sigma - \mathbf{x}_0^i)/\sigma$. These models are trained to predict $\mathbf{x}_0$ directly, following the original EDM and PFGM++ conventions.

Datasets and protocol. We build on the official PFGM++ codebase [10], preserving the original architecture and training configuration as faithfully as possible; the only notable modification is a training batch size of $B = 256$ instead of 512 for CIFAR-10 [24]. All image experiments are

Table 5: FID (↓) [26] on CIFAR-10 as a function of the number of samples N used for the multi-sample target. Best scores per column are **bolded**. Standard deviations are measured across 3 different sets of seeds for FID evaluation.

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

Table 6: FID (↓) at the final checkpoint on FFHQ-64 as a function of the multi-sample target size N. Best scores per column are **bolded**. Standard deviations are measured as in Table 5.

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

unconditional generation on CIFAR-10 and FFHQ-64 [25], trained for 100k kimg on CIFAR-10 and 200k kimg on FFHQ, where one kimg is one thousand training images processed.

For the multi-sample target 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 samples drawn at random from the dataset, so the estimate always uses a pool at least as large as the gradient-evaluation batch.

For each FID [26] evaluation we draw 50k samples and compare against the full training set, with 50 solver steps and the Euler sampler for the two-sided interpolant and the Heun sampler for EDM and PFGM++. All image experiments ran on 2×A100 GPUs; each CIFAR-10 run took about two days and each FFHQ run about eight.

E.2 Results

Table 5 reports CIFAR-10 FID as a function of the number of samples N used for target estimation, and Figure 5 the corresponding training curves.

For the two-sided interpolant nothing changes: the curves are nearly indistinguishable across all tested N (Figure 5a), and the FID stays within 3.03–3.05, its best value reached at $N = 1$.

For EDM and PFGM++ the multi-sample target helps, improving both convergence speed and final FID (Figures 5b–d), though not monotonically in N: EDM is best at $N = 2048$, $2.29 \to 2.12$, while both PFGM++ variants peak at $N = 256$, $2.40 \to 2.24$ for $d = 128$ and $2.40 \to 2.14$ for $d = 2048$. For the trained checkpoints, the observed FID gaps exceed the standard deviations across evaluation seeds (0.02–0.05).

On FFHQ-64 the ordering is the same but the margin is narrower (Figure 6, Table 6): both methods are best at $N = 256$, $2.62 \to 2.51$ for EDM and $2.69 \to 2.56$ for PFGM++ with $d = 128$, and by $N = 2048$ EDM is back at the single-sample level while PFGM++ retains only a small part of the gain. Appendix E.3 ties this narrowing to the resolution.

E.3 Estimator degeneracy at higher resolution

The estimator (42) is a weighted average, so it is only as multi-sample as its weights allow. We measure their concentration by the Gini index of the weights ω_i,

$$G = 1 - \sum_{i=1}^{N} \omega_i^2, \tag{44}$$

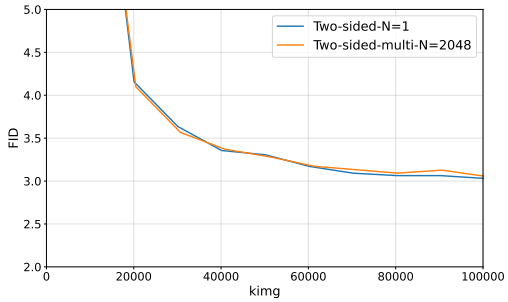

(a) Two-sided interpolant ($s = 0.1$).

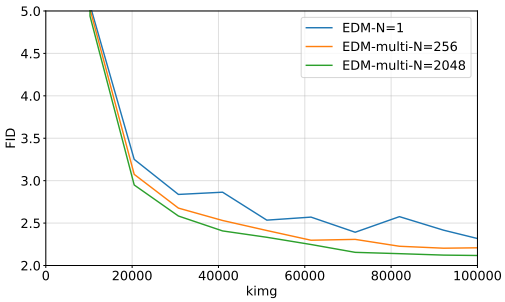

(b) EDM.

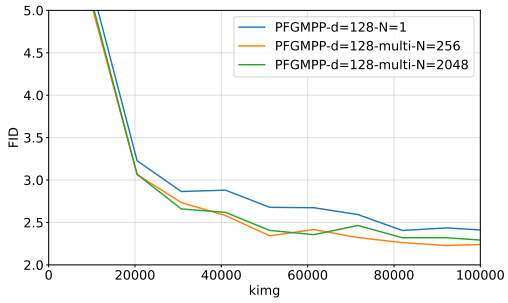

(c) PFGM++ ($d = 128$).

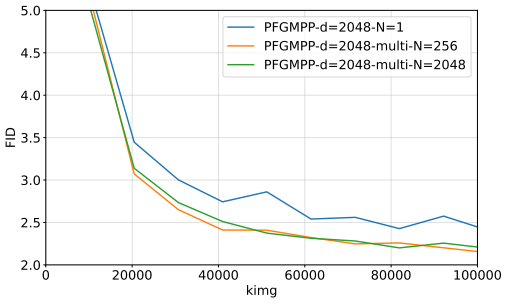

(d) PFGM++ ($d = 2048$).

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

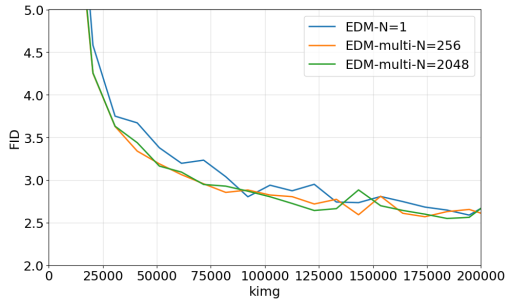

(a) EDM.

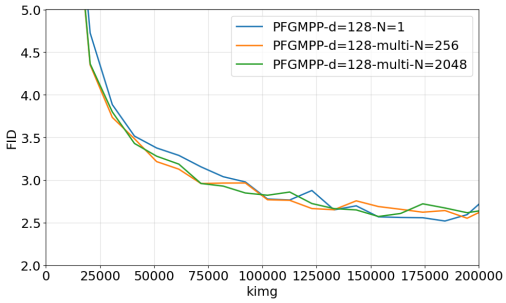

(b) PFGM++ ($d = 128$).

Figure 6: Training convergence on FFHQ-64 (FID vs. kimg) for the single-sample target and multi-sample targets with $N = 256$ and $N = 2048$.

which is zero when one weight carries everything and grows as the weight mass spreads out.

For the two-sided interpolant the weights are degenerate already at 32×32: a single conditional term dominates the sum, and empirically that term is almost always the pair $(\mathbf{x}_0, \mathbf{x}_T)$ used to generate the intermediate point $\mathbf{x}_t \sim p_t^{\mathbf{x}_0,\mathbf{x}_T}$, so additional samples contribute negligibly. Removing that pair from the estimator drives the FID above 100, confirming that the remaining samples cannot replace it. For EDM and PFGM++ the weights are less concentrated and spread further as N grows, which is where the gains of Appendix E.2 come from. This also suggests that the estimator could be improved by relevance-based sampling of conditional pairs instead of uniform sampling.

Table 7 shows how the concentration depends on dimensionality: as the image resolution grows, G falls by orders of magnitude, so the multi-sample target drifts towards the single-sample one.

Table 7: Gini index G of the weights of the multi-sample target (42) for the EDM conditional path $p_\sigma(\mathbf{x}_\sigma \mid \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_0, \sigma^2 I)$, with σ drawn from the noise distribution proposed by EDM. Columns vary the resolution of the CIFAR images the target samples are drawn from; rows vary the target-set size N.

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

E.4 Conclusion

The IFM estimator transfers: for the one-sided methods a multi-sample target improves both convergence and final FID. Its contribution, however, shrinks as the dimensionality grows, the gain at FFHQ 64×64 being noticeably smaller than at CIFAR-10 32×32, and this is exactly what the weight concentration predicts: in high dimension one conditional term dominates, and the multi-sample estimator degenerates into the single-sample one.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes].

   Justification: The main claims are supported by Sections 3–4 and the empirical validations in Appendix D.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes].

   Justification: The theoretical scope and assumptions are stated in Sections 3–4 and Appendices B.2–B.4. Practical limitations are discussed in Section 2.2 and Appendix E.3.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes].

Justification: Assumptions and proofs are provided in Appendices B and C. The orientation argument underlying Corollary 4.3 is given in Section 4.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes].

Justification: Experimental protocols are described in Appendices D and E; architectures, optimizers and remaining hyperparameters are provided in the public code at `https://github.com/daniil-shlenskii/flow-field-matching-duality`.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes].

   Justification: The code is publicly released at `https://github.com/daniil-shlenskii/flow-field-matching-duality`.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes].

   Justification: Training and evaluation settings are described in Appendices D.1 and E.1, with architectures, optimizers and remaining hyperparameters in the released code.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: For each trained model, we report mean $\pm$ standard deviation over three evaluations: FID with different evaluation seed sets (Tables 5 and 6) and W_2 with resampled source subsets (Table 4).

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes].

   Justification: Compute resources and training times for the image experiments are reported in Appendix E.1.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes].

   Justification: The work is theoretical and uses public benchmark data (Appendices D.1 and E.1), with no human-subject studies or consequential deployment.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A].

    Justification: The work establishes a theoretical relation between existing frameworks and introduces no new model, dataset or deployment; we identify no societal impact specific to this contribution.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A].

    Justification: We release experimental code, with no high-risk pretrained model or dataset requiring special safeguards.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes].

    Justification: CIFAR-10 [24] is cited as requested and has no explicit license. FFHQ [25] and PFGM++ [10] use CC BY-NC-SA 4.0; individual FFHQ images retain their CC BY 2.0, CC BY-NC 2.0, Public Domain Mark 1.0, CC0 1.0 or U.S. Government Works licenses. All assets are used for non-commercial research.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.

- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `https://paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes].

    Justification: The released code is documented in the README at `https://github.com/daniil-shlenskii/flow-field-matching-duality`.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A].

    Justification: The work involves no crowdsourcing or human-subject studies.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A].

    Justification: The work involves no human-subject studies requiring IRB review.

    Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: LLMs are not a component of the proposed methods or experimental pipeline.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.