Just Add More Capacitors: Eliminating Flux Leakage in Electrostatic Field Matching

Anonymous Authors[1]

Abstract

Electrostatic field-line methods provide a physics-inspired view of generative modeling, originating in noise-to-data generation with Poisson Flow Generative Models. Electrostatic Field Matching (EFM) extends this idea to data-to-data generative modeling by placing source and target distributions on oppositely charged plates in an augmented space. However, the single-capacitor geometry of EFM induces backward-oriented and exterior field lines, making the induced generative transport stochastic and complicating practical ODE-based inference. To this end, we propose Capacitor Chain Field Matching (C^{2}FM), a generative modeling framework for data-to-data transport based on a periodic chain of electrostatic capacitors. Instead of using a single pair of charged plates, C^{2}FM places infinitely many alternating copies of the source and target distributions along the auxiliary dimension. This periodic construction creates a symmetry that eliminates flux leakage beyond adjacent source and target plates, thereby preventing backward-oriented trajectories. Consequently, samples can be transported deterministically by integrating a forward-oriented field from the source distribution to the target distribution. We prove that the induced field-line map transports the source distribution to the target distribution and provide a practical field-matching algorithm. Our framework preserves the electrostatic interpretation of EFM while yielding a geometrically well-behaved generative transport mechanism for data-to-data transfer.

1. Introduction

Diffusion models (Song et al., 2020; Karras et al., 2022), Flow Matching (Lipman et al., 2022; Liu et al., 2022), and Stochastic Interpolants (Albergo & Vanden-Eijnden, 2023; Albergo et al., 2025) have become standard tools for both noise-to-data generation and data-to-data transport. These methods typically learn a velocity field whose integral curves transform one distribution into another.

A complementary line of work proposes to construct similar generative fields from electrostatic principles (Xu et al., 2023; Shlenskii & Korotin, 2026). Poisson Flow Generative Models (Xu et al., 2022, PFGM) interpret the data distribution as an electric charge distribution embedded in an augmented space. The corresponding electrostatic field, obtained by the superposition principle, defines field lines that connect a simple prior distribution to the data distribution. This provides a physics-inspired mechanism for noise-to-data generation, in which samples are generated by following electrostatic field lines.

Electrostatic Field Matching (Kolesov et al., 2025, EFM) extends this idea to data-to-data transport. Given two distributions, EFM places them on two parallel hyperplanes in an augmented space and assigns them opposite electric charges. The resulting electrostatic field induces field lines between the two charged plates, thereby defining a transport from the source to the target distribution. This construction is theoretically appealing because the electrostatic field is determined solely by the source and target marginals through superposition, rather than by a prescribed coupling between samples.

However, the single-capacitor geometry of EFM also introduces several practical difficulties (Figure 1a). Field lines are not confined to the slab between the two plates: some lines leave the region, some are backward-oriented, and some cross beyond the target plate before returning. As a result, exact EFM transport requires stochastic decisions over field-line branches, while practical ODE-based inference can become unstable when the vertical component of the field changes sign. These issues make the original EFM construction difficult to train and integrate in practice.

Interaction Field Matching (Manukhov et al., 2026, IFM)

---

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

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

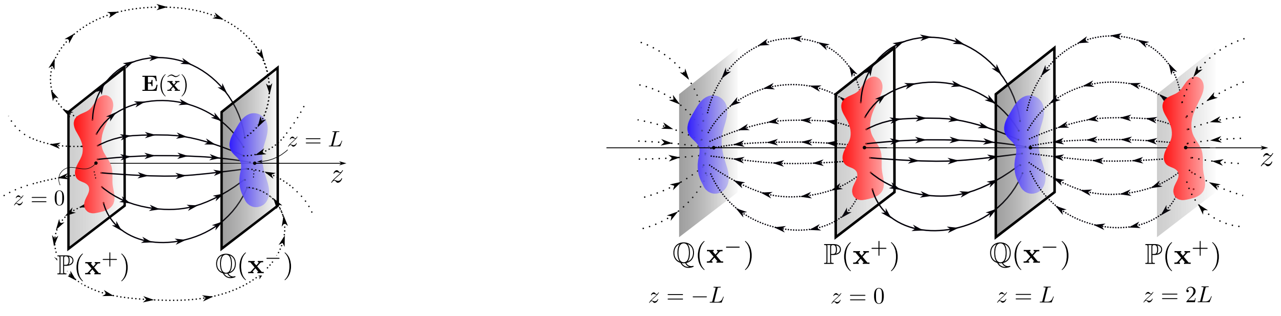

(a) EFM (Kolesov et al., 2025)

(b) Our C^2FM

Figure 1. Electrostatic Field Matching (Kolesov et al., 2025, EFM) and our Chain Capacitor Field Matching (C^2FM) framework. **(a)** In EFM, two D-dimensional data distributions $\mathbb{P}(\cdot)$ and $\mathbb{Q}(\cdot)$ are embedded in $\mathbb{R}^{D+1}$ at $z = 0$ and $z = L$, respectively. The first distribution is assigned a positive charge, the second a negative charge, creating a capacitor-like electric field. Movement along the resulting field lines transports one distribution into the another one. Nevertheless, in EFM some field lines leave the region $z \in [0, L]$: both backward- and forward-oriented lines arise, and a correct transport must account for both families. As a consequence, EFM faces difficulties in training and inference (see §2.2 for details). **(b)** In C^2FM, we introduce infinite number of the positive and negative distribution copies. $\mathbb{P}(\cdot)$ is placed on the planes $z = 0, \pm 2L, \pm 4L, \dots$, while $\mathbb{Q}(\cdot)$ is placed on $z = \pm L, \pm 3L, \pm 5L, \dots$, forming a periodic chain of capacitors. This charge configuration is symmetric under reflection across every plane $z = nL,\ n \in \mathbb{Z}$, which resolves the principal limitations of EFM while staying fully within the electrostatic paradigm (see §3.2).

was proposed as a general field-based framework that contains EFM as a special case. Within this framework, one can choose interaction fields with more convenient field-line geometry, thereby avoiding some of the practical difficulties of EFM. However, recent work (Shlenskii et al., 2026) has shown that certain practical IFM constructions coincide with existing flow-based stochastic interpolants, and thus no longer preserve the electrostatic structure of EFM. This raises a question: can one overcome the practical limitations of EFM while preserving its electrostatic nature?

In this work, we answer this question affirmatively. We propose $\mathrm{C}^2\mathrm{FM}$ (Figure 1b), a capacitor-chain field matching method for data-to-data generative modeling. Instead of placing the source and target distributions on a single pair of charged plates, we place infinitely many alternating copies of the two distributions along the auxiliary dimension. Positive copies of the source distribution are placed at $z = 2nL$, while negative copies of the target distribution are placed at $z = (2n+1)L$, for all $n \in \mathbb{Z}$. This periodic electrostatic construction creates a reflection symmetry that confines forward field lines to the slab between adjacent plates. As a result, $\mathrm{C}^2\mathrm{FM}$ yields a practical generative algorithm that overcomes the main limitations of EFM while preserving its electrostatic interpretation.

Our **contributions** are as follows:

- We introduce $\mathrm{C}^2\mathrm{FM}$, a chain of electrostatic capacitors that preserves the electrostatic interpretation of EFM while confining field lines between adjacent source and target plates. We prove that the induced field-line map deterministically transports the source to the target distribution (§3, §3.2).
- We derive a practically applicable field-matching training objective and an ODE-based inference. We empirically demonstrate that $\mathrm{C}^2\mathrm{FM}$ improves sample quality and trajectory stability compared with the original EFM (§3.4, §4.1).

2. Background

This section reviews the electrostatic concepts used throughout the paper. We first recall the superposition principle and field lines in §2.1, then summarize EFM and its limitations in §2.2.

2.1. Electrostatics

We start with the fundamental principles of electrostatics necessary for understanding electrostatic-based generative models. A detailed treatment of three-dimensional electrostatics can be found in any standard electricity and magnetism textbook, e.g., (Landau & Lifshitz, 1971, Chapter 5). The generalization of electrostatics to high-dimensional spaces is discussed in (Caruso et al., 2023) and (Kolesov et al., 2025, Appendix A).

The electrostatic field. Let $q : \mathbb{R}^D \to \mathbb{R}$ be the density of a charge distribution on $\mathbb{R}^D$. The distribution may contain both positive and negative charges and is assumed to have finite total charge ($\int |q(\mathbf{x})| d\mathbf{x} < \infty$). At a point $\mathbf{x} \in \mathbb{R}^D$ it produces the electrostatic field $\mathbf{E} : \mathbb{R}^D \to \mathbb{R}^D$:

$$\mathbf{E}(\mathbf{x}) = \int \frac{1}{S_{D-1}} \frac{(\mathbf{x} - \mathbf{x}')}{||\mathbf{x} - \mathbf{x}'||^D} q(\mathbf{x}') d\mathbf{x}', \qquad (1)$$

where S_{D-1} is the surface area of an $(D-1)$-dimensional sphere with unit radius. Equation (1) represents the field as a superposition of point-charge fields.

Electric field strength lines. An electric field strength line is a curve $\mathbf{x}(\tau) \in \mathbb{R}^D,\ \tau \in [a,b] \subset \mathbb{R}$ whose tangent to each point is parallel to the electric field at that point. In other words:

$$\frac{d\mathbf{x}(\tau)}{d\tau} = \mathbf{E}(\mathbf{x}). \tag{2}$$

Electric field lines are a key concept for electrostatic generative models such as PFGM and EFM.

2.2. Electrostatic Field Matching (EFM)

The first generative models based on electrostatics, Poisson Flow Generative Models (PFGM and PFGM++) (Xu et al., 2022; 2023), addressed the problem of mapping a simple noise distribution to a complex data distribution. To handle data-to-data transport, Electrostatic Field Matching (EFM) (Kolesov et al., 2025) extends the idea by introducing charge distributions of opposite signs. In EFM, two data distributions $\mathbb{P}(\mathbf{x}^+)$ and $\mathbb{Q}(\mathbf{x}^-)$ on $\mathbb{R}^D$ are embedded into an extended space $\mathbb{R}^{D+1}$ as charged hyperplanes: positive copies $\mathbb{P}$ are placed on the plane $z = 0$ and negative copies $\mathbb{Q}$ on $z = L$. This arrangement mimics a $(D+1)$-dimensional **capacitor**.

Points in the extended space are written as $\widetilde{\mathbf{x}} = (\mathbf{x}, z) \in \mathbb{R}^{D+1}$. The total electric field is the superposition of the fields generated by the two plates:

$$\mathbf{E}(\widetilde{\mathbf{x}}) = \mathbf{E}^+(\widetilde{\mathbf{x}}) + \mathbf{E}^-(\widetilde{\mathbf{x}}), \tag{3}$$

where $\mathbf{E}^\pm$ are obtained from the charge densities $\mathbb{P}$ and $\mathbb{Q}$ via the high-dimensional Coulomb integral (1) in $D+1$ dimensions.

The key property proved in (Kolesov et al., 2025) is that the integral curves of this field, i.e. the solutions of $\frac{\mathrm{d}\widetilde{\mathbf{x}}}{\mathrm{d}\tau} = \mathbf{E}(\widetilde{\mathbf{x}})$, almost surely transfer a sample of $\mathbb{P}$ (on $z = 0$) to a sample of $\mathbb{Q}$ (on $z = L$), but there are several tricky nuances which we discuss in the limitations section below.

To approximate this transport in practice, one trains a neural network $f_\theta : \mathbb{R}^{D+1} \to \mathbb{R}^{D+1}$ to predict the normalized field direction. The target is constructed by estimating $\mathbf{E}$ through the superposition principle (3), where the integrals are replaced by Monte Carlo sums over the data. The network is optimized with the regression loss

$$\mathbb{E}_{\widetilde{\mathbf{x}}} \left\| f_\theta(\widetilde{\mathbf{x}}) - \frac{\mathbf{E}(\widetilde{\mathbf{x}})}{\|\mathbf{E}(\widetilde{\mathbf{x}})\|} \right\|^2 \to \min_\theta, \tag{4}$$

where the expectation $\mathbb{E}_{\widetilde{\mathbf{x}}}$ is taken over points in the region between the two plates. This set of points is called the training volume.

Limitations of EFM. While EFM formally establishes a transport between the two distributions, it has a few drawbacks that arise from the behaviour of electric field lines outside the region $[0, L]$:

1. **Field lines leaving $[0, L]$.** The field lines emanating from the positive plate split into two families: forward-oriented ($z > 0$) and backward-oriented ($z < 0$), see Figure 1a. Both families can exit the interval $[0, L]$: forward lines may overshoot into $z > L$, while backward lines go into $z < 0$. Consequently, any **training** objective must cover a larger volume, extending into $z < 0$ and $z > L$, which complicates the regression task.
2. **Stochasticity of the transport map.** To fully cover the target distribution, both forward and backward families of field lines must be taken into account; using only forward-oriented lines leads to an incomplete mapping. Consequently, the exact transport map from $\mathbb{P}$ to $\mathbb{Q}$ is not deterministic. First, at a starting point one must randomly choose between the forward and the backward line with probability proportional to the outgoing flux (Kolesov et al., 2025, §3.2). Second, not all forward-oriented lines terminate directly at $z = L$; some of them leave the region $[0, L]$ before returning. For those lines, a fraction of the probability is deposited at $z = L$ on first approach, while the remainder continues into $z > L$, requiring a second stochastic decision.
3. **Inference difficulties from lines crossing $z > L$.** Standard EFM implementations (Kolesov et al., 2025, Algorithms 1, 2) use only forward-oriented lines which do not guarantee valid transport. For the subset of these lines that cross into $z > L$, at some point the field turns back toward $z = L$; at that turning point $\mathbf{E}_z$ changes sign, leading to numerical instability during ODE integration and the correct inference even for the forward-oriented lines hard. Thus, the integration heuristically stops at $z = L$.

3. Capacitor Chain Field Matching (C^{2}FM)

This section presents the proposed C^{2}FM construction. We first give the geometric intuition in §3.1, then state the main theoretical construction in §3.2, relate the method to prior field-based models in §3.3, and describe learning and inference in §3.4.

3.1. Intuitive Explanation of the Method

To address the limitations of EFM, we propose a new charge arrangement geometry. Unlike EFM, where the distributions $\mathbb{P}(\widetilde{\mathbf{x}})$ and $\mathbb{Q}(\widetilde{\mathbf{x}})$ are placed on the planes $z = 0$ and $z = L$ respectively, we propose creating an infinite number of copies of $\mathbb{P}(\widetilde{\mathbf{x}})$ and $\mathbb{Q}(\widetilde{\mathbf{x}})$. The positive charge distributions are placed on the planes $z = 0, \pm 2L, \pm 4L, \dots$ while the negative ones on $z = \pm L, \pm 3L, \pm 5L, \dots$ (see Fig.1). This structure is invariant under reflection with respect to any plane $z = nL$, $n \in \mathbb{Z}$, which ensures that the z-projection of the field on these planes is exactly zero. Consequently, the electric field in the region $z \in (0, L)$ is confined strictly between the planes $z = 0$ and $z = L$: no field line that starts in $z \in (0, L)$ ever leaves $z \in (0, L)$. This resolves both the problem of field lines leaving the interval $(0, L)$ and the stochasticity of the transport map. Moreover, this configuration guarantees $\mathbf{E}_z > 0$ throughout $(0, L)$, which stabilizes inference.

3.2. Main Theoretical Construction

Let $\mathbb{P}(\mathbf{x})$ and $\mathbb{Q}(\mathbf{x})$, with $\mathbf{x} \in \mathbb{R}^D$, be two data distributions. We assign to the first distribution a positive charge and to the second distribution a negative charge. We place these distributions in the extended space $\mathbb{R}^{D+1}$ as an infinite periodic array, as described in §3.1. The charge distributions can be written with Dirac delta functions as:

$$
\begin{aligned}
q^+(\widetilde{\mathbf{x}}) &= q^+(\mathbf{x}, z) = \mathbb{P}(\mathbf{x}) \sum_{n=-\infty}^{\infty} \delta(z - 2nL), \\
q^-(\widetilde{\mathbf{x}}) &= q^-(\mathbf{x}, z) = -\mathbb{Q}(\mathbf{x}) \sum_{n=-\infty}^{\infty} \delta(z - (2n+1)L).
\end{aligned} \tag{5}
$$

Using these definitions, the electric field $\mathbf{E}^+(\widetilde{\mathbf{x}})$ generated by the positive charges and $\mathbf{E}^-(\widetilde{\mathbf{x}})$ generated by the negative charges are obtained via the superposition principle in $(D+1)$-dimensional point $\widetilde{\mathbf{x}} = (\mathbf{x}, z)$ (see (1)):

$$
\begin{aligned}
&\mathbf{E}^\pm(\widetilde{\mathbf{x}}) = \int \frac{\widetilde{\mathbf{x}} - \widetilde{\mathbf{x}}'}{S_D \|\widetilde{\mathbf{x}} - \widetilde{\mathbf{x}}'\|^{D+1}} \, q^\pm(\widetilde{\mathbf{x}}') \, d^{D+1}\widetilde{\mathbf{x}}' \iff \\
&\mathbf{E}^+(\widetilde{\mathbf{x}}) = \sum_{n=-\infty}^{\infty} \int \frac{\big(\mathbf{x} - \mathbf{x}', \; z - 2nL\big)\mathbb{P}(\mathbf{x}') \, d^D\mathbf{x}'}{S^D\big(\|\mathbf{x} - \mathbf{x}'\|^2 + (z - 2nL)^2\big)^{\frac{D+1}{2}}}, \\
&\mathbf{E}^-(\widetilde{\mathbf{x}}) = \sum_{n=-\infty}^{\infty} \int \frac{-\big(\mathbf{x} - \mathbf{x}', \; z - (2n+1)L\big)\mathbb{Q}(\mathbf{x}') \, d^D\mathbf{x}'}{S^D\big(\|\mathbf{x} - \mathbf{x}'\|^2 + (z - (2n+1)L)^2\big)^{\frac{D+1}{2}}}
\end{aligned} \tag{6}
$$

where S_D is the surface area of a D-dimensional unit sphere. The total electric field is then

$$
\mathbf{E}(\widetilde{\mathbf{x}}) = \mathbf{E}^+(\widetilde{\mathbf{x}}) + \mathbf{E}^-(\widetilde{\mathbf{x}}). \tag{7}
$$

Theorem 3.1 (Only forward-oriented field lines)**.** *Let a forward-oriented field line* $d\widetilde{\mathbf{x}} = \mathbf{E}(\widetilde{\mathbf{x}})d\tau$ *start at a point* $\widetilde{\mathbf{x}} = (\mathbf{x}, z)$ *with* $\mathbf{x} \in \operatorname{supp}\mathbb{P}$[1], $z \to 0^+$, *where* $\mathbf{E}(\widetilde{\mathbf{x}})$ *is the total electric field from (7). Then this field line never leaves the boundaries* $z \in [0, L]$ *and always satisfies* $\mathbf{E}_z > 0$ *on* $z \in (0, L)$.

This statement guarantees that the field configuration (7) resolves problems 1 and 3, observed in EFM (see §2.2). The proof of the theorem is given in §A.

Finally, let us define the map between the *deterministic* distributions $T : \operatorname{supp}(\mathbb{P}(\mathbf{x}^+)) \to \operatorname{supp}(\mathbb{Q}(\mathbf{x}^-))$ using electric field lines (7). Let us consider the field (7) exclusively between the planes $z = 0$ and $z = L$. The region $z \in [0, L]$ is sufficient because Theorem 3.1 guarantees that field lines starting in the region $z \in [0, L]$ never leave it.

Consider a point $\widetilde{\mathbf{x}}^+ = (\mathbf{x}^+, 0)$ in the support of the first distribution (positive plate). Let us denote the field line $\widetilde{\mathbf{x}}(t), (t \in [a, b])$ starting at this point in forward direction. From the properties of electric field lines (Kolesov et al., 2025, Appendix A) and the confinement theorem (Theorem 3.1), it must end almost surely at the point $\widetilde{\mathbf{x}}^- = (\mathbf{x}^-, L)$ in the support of the second distribution $\mathbb{Q}(\widetilde{\mathbf{x}}^-)$ (negative plate). Thus, $\widetilde{\mathbf{x}}(a) = (\mathbf{x}^+, 0) = \widetilde{\mathbf{x}}^+, \;\; \widetilde{\mathbf{x}}(b) = (\mathbf{x}^-, L) = \widetilde{\mathbf{x}}^-$. Using this, we define $T(\mathbf{x}^+) = \mathbf{x}^-$.

Theorem 3.2 (**C^{2}FM transport theorem**)**.** *Let* $\mathbb{P}(\mathbf{x}^+)$ *and* $\mathbb{Q}(\mathbf{x}^-)$ *be two data distributions with compact support, and let* T *be the deterministic map defined above. Let* $\mathbf{x}^+$ *be distributed as* $\mathbb{P}(\mathbf{x}^+)$. *Then* $\mathbf{x}^- = T(\mathbf{x}^+)$ *is distributed as* $\mathbb{Q}(\mathbf{x}^-)$ *almost surely:*

$$
\textit{If } \; \mathbf{x}^+ \sim \mathbb{P}(\mathbf{x}^+) \Rightarrow T(\mathbf{x}^+) = \mathbf{x}^- \sim \mathbb{Q}(\mathbf{x}^-). \tag{8}
$$

Thus, the electric field lines in our periodic configuration define a deterministic map T that transfers $\mathbb{P}$ to $\mathbb{Q}$, resolving the stochasticity issue of EFM (see §2.2). We note that although the theorem may look similar to the analogous one in EFM, the difference is crucial: in EFM the map is stochastic, while in our case it is deterministic. The proof of Theorem 3.2 is given in §A.

3.3. Relation to Other Methods

C^{2}FM is an electrostatic generative model. The first works in this direction, (Xu et al., 2022, PFGM) and (Xu et al., 2023, PFGM++), generalize diffusion models, but like early diffusion models they are limited to noise-to-data generation. To handle data-to-data transfer, (Kolesov et al., 2025, EFM) was introduced; the method itself and its limitations are discussed in §2.2.

A subsequent approach, (Manukhov et al., 2026, Interaction Field Matching (IFM)), addresses the problems of EFM by broadening the electrostatic perspective to the general framework of *interaction fields*. In this framework, arbitrary fields act between pairs of "charged" particles. As long as the field satisfies a few simple conditions — field lines start and end at the charges, the flux is conserved, and a generalized superposition principle holds — it provably transports one data distribution to the other. Many concrete realizations of IFM are possible. The authors of (Manukhov

---

[1] $\operatorname{supp}(f) = \{x \in X | f(x) \neq 0\}$, where X is the domain of function f

et al., 2026) propose a field inspired by the strong interaction of subnuclear particles (quarks) in physics. This realization resolves the EFM issues discussed in §2.2.

From this perspective, C^{2}FM, along with PFGM and EFM, constitutes a particular realization of the IFM framework. However, C^{2}FM differs fundamentally from the strong-interaction-based realization of IFM in two respects. First, the IFM realization introduces the field of a *pair* of charges, $\mathbf{E}_{q\bar{q}}$, which cannot be decomposed as $\mathbf{E}_q + \mathbf{E}_{\bar{q}}$, whereas in electrostatics the superposition principle (7) holds. Second, the strong-interaction realization exhibits an explicit dependence on the transport plan $\pi(\cdot,\cdot)$, while our method is transport-plan agnostic.

For further context, we mention the work of (Shlenskii et al., 2026), which uncovers a deep link between IFM and Conditional Flow Matching (CFM). It demonstrates that conventional CFM is a subset of IFM, termed *forward-only-IFM*, and that other IFM realizations, in particular EFM, are strictly more expressive than CFM.

3.4. Learning and Inference Algorithm

To move between data distributions, it is sufficient to follow the electric field lines. In our periodic configuration, the field $\mathbf{E}(\widetilde{\mathbf{x}})$ defined in (7) is confined to the slab $z \in [0, L]$ for all forward-oriented lines starting from the positive plate. Theorem 3.1 guarantees that such lines never leave this slab and always satisfy $\mathbf{E}_z > 0$, which eliminates the problems of backward-oriented lines and lines crossing $z > L$ that plagued the original EFM.

Training. To recover the electric field $\mathbf{E}(\cdot)$ for points in $(D+1)$-dimensional space between the hyperplanes, we approximate it with a neural network $f_\theta(\cdot) : \mathbb{R}^{D+1} \to \mathbb{R}^{D+1}$. The approximation requires a training volume. We sample a value t from the uniform distribution $\mathcal{U}(0, L)$ and take two random samples $\widetilde{\mathbf{x}}^+$ (from the positive plate at $z = 0$) and $\widetilde{\mathbf{x}}^-$ (from the negative plate at $z = L$). We then generate an intermediate point $\widetilde{\mathbf{x}}$ via linear interpolation and noise:

$$\widetilde{\mathbf{x}} = \frac{t}{L}\widetilde{\mathbf{x}}^- + \left(1 - \frac{t}{L}\right)\widetilde{\mathbf{x}}^+ + \widetilde{\varepsilon}(z), \tag{9}$$

where $\widetilde{\varepsilon}(z) = 0$ in 2D experiments (§4.1). In translation experiments (§4.2), we sample z from uniform distribution on $[0, L]$ and $\widetilde{\epsilon}(z)$ from $\mathcal{N}(0, \sigma^2(z)I)$, where the variance $\sigma^2(z)$ is $= \frac{L}{2} - |\frac{L}{2} - z|$. The noise term enlarges the training volume and improves generalization.

The ground-truth $\mathbf{E}(\widetilde{\mathbf{x}})$ is estimated using (7). In practice, the infinite sums over n are truncated to a finite number of capacitor copies C (e.g., $|n| \leq C$) because the contributions decay rapidly with distance. The integrals are approximated by Monte Carlo sampling from $\mathbb{P}(\mathbf{x}^+)$ and $\mathbb{Q}(\mathbf{x}^-)$. We then train the neural network $f_\theta(\widetilde{\mathbf{x}})$ to learn the *normalized* electric field $\frac{\mathbf{E}(\widetilde{\mathbf{x}})}{\|\mathbf{E}(\widetilde{\mathbf{x}})\|}$, as this has been found to work better than learning the unnormalized field. The learning objective is:

$$\mathbb{E}_{\widetilde{\mathbf{x}}} \left\| f_\theta(\widetilde{\mathbf{x}}) - \frac{\mathbf{E}(\widetilde{\mathbf{x}})}{\|\mathbf{E}(\widetilde{\mathbf{x}})\|} \right\|^2 \to \min_\theta . \tag{10}$$

Inference. Having learned the normalized vector field $f_\theta(\cdot) \approx \frac{\mathbf{E}(\cdot)}{\|\mathbf{E}(\cdot)\|}$, we simulate the movement between hyperplanes to transfer data from $\mathbb{P}(\mathbf{x}^+)$ to $\mathbb{Q}(\mathbf{x}^-)$. Because $\mathbf{E}_z > 0$ along the entire trajectory (Theorem 3.1), we can reparameterize the ODE using z as the integration variable:

$$\begin{aligned} d\widetilde{\mathbf{x}} = d(\mathbf{x}, z) &= \left(\frac{d\mathbf{x}}{dt}\frac{dt}{dz}dz,\ dz\right) = \big(\mathbf{E}_x(\widetilde{\mathbf{x}})\,\mathbf{E}_z^{-1}(\widetilde{\mathbf{x}}),\ 1\big)\,dz \\ &= \left(\frac{\mathbf{E}_x(\widetilde{\mathbf{x}})}{\|\mathbf{E}(\widetilde{\mathbf{x}})\|}\,\frac{\|\mathbf{E}(\widetilde{\mathbf{x}})\|}{\mathbf{E}_z(\widetilde{\mathbf{x}})},\ 1\right) dz \approx \big(f_\theta(\widetilde{\mathbf{x}})_x\, f_\theta(\widetilde{\mathbf{x}})_z^{-1},\ 1\big)\,dz, \end{aligned} \tag{11}$$

where $f_\theta(\widetilde{\mathbf{x}}) = (f_\theta(\widetilde{\mathbf{x}})_x, f_\theta(\widetilde{\mathbf{x}})_z)$. We start with samples from $\mathbb{P}(\mathbf{x}^+)$ at $z = 0$ and integrate forward until z reaches L. Since the field line never leaves $[0, L]$ and $\mathbf{E}_z > 0$, the integration is well-defined and no special stopping criterion is needed. This resolves the inference instability of original EFM.

Training steps are summarized in Algorithm 1. Inference algorithm coincides with (**?**)Algorithm 2]kolesov2025field.

4. Experimental Illustrations

We present proof-of-concept experiments for our C^2FM method. First, we conduct a 2D illustrative experiment (§4.1) and provide a study of the sensitivity of our model to the choice of the number of capacitors. Second, we perform image-to-image translation experiments (§4.2), where we transfer colored digits '2' to colored digits '3'. Additionally, we provide the technical implementation details, memory usage and inference time in Appendix B.

4.1. 2D Experiments

We evaluate the proposed **C^{2}FM** method on synthetic two-dimensional datasets, where the geometry of the learned transport can be inspected directly. The main goal of these experiments is to study how the truncated periodic capacitor construction affects the quality of generative transport and to verify that increasing the number of capacitor copies improves the learned field and the resulting samples.

In all experiments, the source and target distributions are defined on $\mathbb{R}^2$ as Gaussian mixtures and are embedded into the augmented space $\mathbb{R}^3$ by introducing the auxiliary coordinate z. The source distribution is placed on the plane $z = 0$, while the target distribution is placed on the plane $z = L$. We fix the plate distance to $L = 20$ in all reported experiments.

Algorithm 1 C^{2}FM Training

```
Input: Distributions accessible by samples:
    $\mathbb{P}(\mathbf{x}^+)\delta(z)$ and $\mathbb{Q}(\mathbf{x}^-)\delta(z-L)$;
    Neural network $f_\theta(\cdot) : \mathbb{R}^{D+1} \to \mathbb{R}^{D+1}$;
    Plate distance $L > 0$; number of periodic copies $C \in \mathbb{N}$;
Output: Learned normalized field $f_\theta(\cdot)$ approximating $\frac{\mathbf{E}(\cdot)}{\|\mathbf{E}(\cdot)\|}$ for $z \in [0, L]$.
Repeat until converged:    Sample a batch $\mathcal{B}^+$ of points $\widetilde{\mathbf{x}}^+ \sim \mathbb{P}(\mathbf{x}^+)\delta(z)$;
    Sample a batch $\mathcal{B}^-$ of points $\widetilde{\mathbf{x}}^- \sim \mathbb{Q}(\mathbf{x}^-)\delta(z-L)$;
    Sample a batch of times $t \sim \mathcal{U}(0, L)$;
    Sample a batch of noise $\widetilde{\varepsilon}$;
    Construct $\widetilde{\mathbf{x}} = \frac{t}{L}\widetilde{\mathbf{x}}^- + (1 - \frac{t}{L})\widetilde{\mathbf{x}}^+ + \widetilde{\varepsilon}$;
    For each $\widetilde{\mathbf{x}}$ in the batch $\mathcal{B}^+$, estimate $\mathbf{E}^+(\widetilde{\mathbf{x}})$ via Monte Carlo:
        $\mathbf{E}^+(\widetilde{\mathbf{x}}) \approx \frac{1}{S_D} \sum_{n=-C}^{C} \frac{1}{|\mathcal{B}^+|} \sum_{\mathbf{x}'_i \in \mathcal{B}^+} \frac{\left(\mathbf{x} - \mathbf{x}'_i,\ z - 2nL\right)}{\left(\|\mathbf{x} - \mathbf{x}'_i\|^2 + (z - 2nL)^2\right)^{(D+1)/2}}$;
    Similarly, estimate $\mathbf{E}^-(\widetilde{\mathbf{x}})$ by averaging over $\mathcal{B}^-$:
        $\mathbf{E}^-(\widetilde{\mathbf{x}}) \approx \frac{1}{S_D} \sum_{n=-C}^{C} \frac{1}{|\mathcal{B}^-|} \sum_{\mathbf{x}'_j \in \mathcal{B}^-} \frac{\left(\mathbf{x} - \mathbf{x}'_j,\ z - (2n+1)L\right)}{\left(\|\mathbf{x} - \mathbf{x}'_j\|^2 + (z - (2n+1)L)^2\right)^{(D+1)/2}}$;
    Compute total field $\mathbf{E}(\widetilde{\mathbf{x}}) = \mathbf{E}^+(\widetilde{\mathbf{x}}) + \mathbf{E}^-(\widetilde{\mathbf{x}})$;
    Compute loss $\mathcal{L} = \mathbb{E}_{\widetilde{\mathbf{x}}} \left\| f_\theta(\widetilde{\mathbf{x}}) - \frac{\mathbf{E}(\widetilde{\mathbf{x}})}{\|\mathbf{E}(\widetilde{\mathbf{x}})\|} \right\|^2 \to \min_\theta$;
    Update $\theta$ using $\frac{\partial \mathcal{L}}{\partial \theta}$;
```

To approximate the ideal infinite periodic chain, we truncate the electrostatic construction and vary the number of capacitor copies C. We compare several values of C, including small truncations and a larger value that more closely approximates the ideal periodic field. We visualize both the generated samples after transport and the corresponding transport trajectories.

We consider two synthetic transport tasks (Figure 2): 5 GMM → 4 GMM and 1 GMM → 2 GMM. These experiments test both mode merging/rearrangement and mode splitting. In both cases, the quality of transport depends strongly on the number of capacitors C. For small C, the generated samples are less well aligned with the target distribution, and in the $5 \to 4$ case the trajectories are less regular.

In particular, for $C = 2$ and, to a lesser extent, for $C = 4$, some trajectories do not terminate on the desired target distribution Q at the corresponding negative plate. Instead, they cross beyond the target plate and enter the exterior region. This region is numerically unstable for inference because it lies outside the training slab where the neural field was fitted, so the learned vector field is less reliable there and may produce inaccurate continuation of the trajectories. These escaping trajectories are highlighted in red in Figure 2. As C increases, this effect is substantially reduced: the field becomes more confined between adjacent plates, the trajectories become smoother and more consistently forward-oriented, and the generated samples better match the target modes. The best qualitative results are obtained for $C = 20$, which most closely approximates the ideal periodic construction. Overall, these experiments show that C^{2}FM can handle both multimodal rearrangement and mode splitting, while consistently benefiting from larger C.

4.2. Handwritten Digits Colored MNIST

We further evaluate C^{2}FM on a higher-dimensional image-to-image translation task using Colored MNIST. In this experiment, the source distribution consists of images of handwritten digit '2', while the target distribution consists of images of handwritten digit '3'. This task tests whether the proposed capacitor-chain construction can be applied beyond low-dimensional synthetic distributions and whether increasing the number of capacitors improves the stability of image-level transport.

Each image is represented as a vector in pixel space and embedded into the augmented space by adding the auxiliary coordinate z. As in the two-dimensional experiments, the source images are placed at $z = 0$ and the target images at $z = L$. We train the neural field to approximate the normalized electrostatic field induced by the source and target empirical distributions, and generate target samples by integrating the learned ODE from the source plane to the target plane.

Figure 3 shows qualitative results for different capacitor trun-

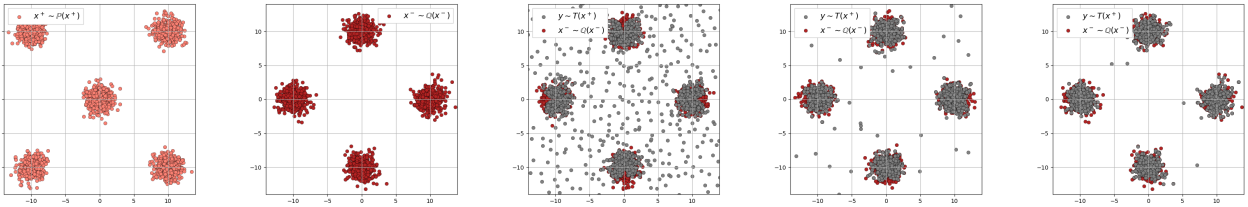

Source (5 GMM) Target (4 GMM) $C = 2$ $C = 4$ $C = 20$

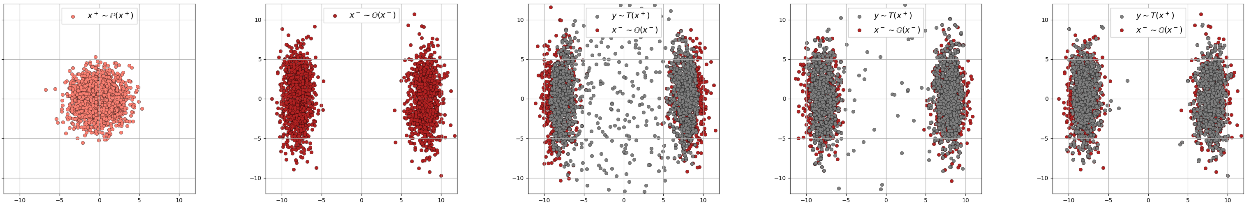

Source (1 GMM) Target (2 GMM) $C = 2$ $C = 4$ $C = 20$

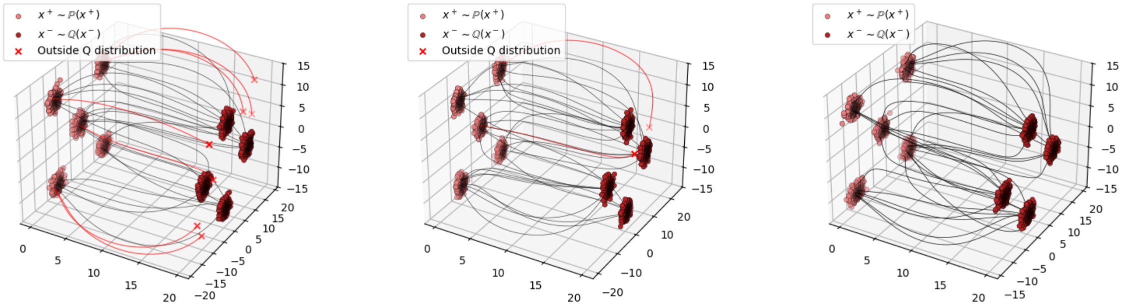

Trajectories, $C = 2$ Trajectories, $C = 4$ Trajectories, $C = 20$

Figure 2. Qualitative results on two-dimensional Gaussian mixture transport tasks. First row: transport from a 5-component GMM to a 4-component GMM. Second row: transport from a single Gaussian to a 2-component GMM. Third row: transport trajectories for the $5 \to 4$ task. Increasing the number of capacitors C improves sample quality and produces more stable trajectories. Trajectories that extend beyond the two planes region are highlighted in red.

cations. With a small number of capacitors, the generated images partially move toward the target digit distribution but may contain unstable or ambiguous digit shapes. In contrast, using a larger truncation, $C = 20$, produces more coherent samples that better resemble the target digit '3'. This behavior is consistent with the two-dimensional experiments: increasing C gives a better approximation to the ideal periodic construction, improves confinement of the learned field, and leads to more stable transport.

5. Discussion and Limitations

We proposed $\mathrm{C}^2\mathrm{FM}$, a capacitor-chain field matching framework for data-to-data generative modeling that preserves the electrostatic structure as EFM while resolving its main geometric difficulties: flux leakage, backward-oriented field lines, and unstable ODE-based inference.

Several limitations still remain. First, some design choices shared with EFM are still heuristic: it is not yet clear how to choose the training volume, or how the plate distance L affects training and inference. Second, the target field computation involves factors of order $\|\mathbf{x} - \mathbf{x}'\|^{-(D+1)}$, which may produce values comparable to machine precision and, as a results, numerical instabilities in high-dimensional settings or at high resolutions. Third, the ideal $\mathrm{C}^2\mathrm{FM}$ construction uses an infinite periodic chain of capacitors, whereas practical implementations must truncate this chain to finitely many copies. This creates a trade-off between field confinement and computational cost.

6. Broader Impacts

This work is primarily foundational research on generative modeling and data-to-data transport. Its potential positive

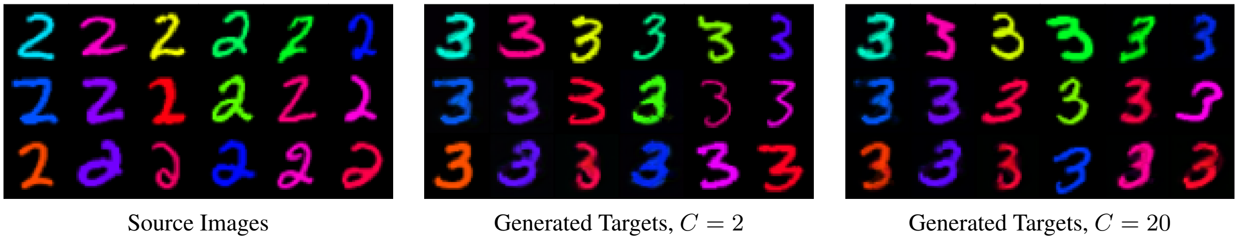

Figure 3. Colored MNIST digit translation with different numbers of capacitors C. Source images are handwritten digits '2', and target images are handwritten digits '3'. Increasing C improves stable transport of generated target samples.

impacts include more stable deterministic transport between distributions, improved understanding of electrostatic generative models, and possible future applications in simulation, domain transfer, and scientific data modeling.

At the same time, improvements in generative modeling methods can contribute to negative downstream uses if applied to sensitive domains, such as generating misleading synthetic media or transforming data in ways that obscure provenance. The present paper evaluates only synthetic low-dimensional distributions and does not release a high-risk pretrained model or scraped dataset. However, future applications of C^2FM to real images, text-conditioned generation, or human-related data should consider privacy, fairness, provenance, and misuse risks, and may require additional safeguards.

References

Albergo, M., Boffi, N. M., and Vanden-Eijnden, E. Stochastic interpolants: A unifying framework for flows and diffusions. *Journal of Machine Learning Research*, 26 (209):1–80, 2025.

Albergo, M. S. and Vanden-Eijnden, E. Building normalizing flows with stochastic interpolants. In *The Eleventh International Conference on Learning Representations*, 2023.

Caruso, F., Oguri, V., and Silveira, F. Still learning about space dimensionality: From the description of hydrogen atom by a generalized wave equation for dimensions d ≥ 3. *American Journal of Physics*, 91(2):153–158, 2023.

Evans, L. C. *Partial Differential Equations*, volume 19 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2nd edition, 2010.

Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

Karras, T., Aittala, M., Aila, T., and Laine, S. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35: 26565–26577, 2022.

Kolesov, A., Manukhov, S., Palyulin, V. V., and Korotin, A. Field matching: an electrostatic paradigm to generate and transfer data. In *Forty-second International Conference on Machine Learning*, 2025. URL `https://openreview.net/forum?id=9dHilxylvC`.

Landau, L. D. and Lifshitz, E. M. *The Classical Theory of Fields (Volume 2 of A Course of Theoretical Physics)*. Pergamon Press, 1971.

Langley, P. Crafting papers on machine learning. In Langley, P. (ed.), *Proceedings of the 17th International Conference on Machine Learning (ICML 2000)*, pp. 1207–1216, Stanford, CA, 2000. Morgan Kaufmann.

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

Lipman, Y., Chen, R. T., Ben-Hamu, H., Nickel, M., and Le, M. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023.

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

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

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

Shlenskii, D., Varlamov, A., Buzun, N., and Korotin, A. Unlocking the duality between flow and field matching. In *ICLR 2026 2nd Workshop on Deep Generative Model in Machine Learning: Theory, Principle and Efficacy*, 2026. URL `https://openreview.net/forum?id=3Vw8H5Scsw`.

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

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

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

A. Proofs of the Theorems

Theorem A.1 (Only forward-oriented field lines)**.** *Let a forward-oriented field line* $d\widetilde{\mathbf{x}} = \mathbf{E}(\widetilde{\mathbf{x}})d\tau$ *start at a point* $\widetilde{\mathbf{x}} = (\mathbf{x}, z)$ *with* $\mathbf{x} \in \operatorname{supp} \mathbb{P}$*,* $z \to 0^+$*, where* $\mathbf{E}(\widetilde{\mathbf{x}})$ *is the total electric field from (7). Then this field line never leaves the boundaries* $z \in [0, L]$ *and always satisfies* $\mathbf{E}_z > 0$ *on* $z \in (0, L)$*.*

Proof. First, we show that for an arbitrary plane $z_0 = n_0 L$ ($n_0 \in \mathbb{Z}$) the field $\mathbf{E}^{\pm}$ possesses the symmetry property

$$\mathbf{E}^{\pm}(\mathbf{x}, z_0 - \epsilon) = \big(\mathbf{E}^{\pm}_x(\mathbf{x}, z_0 + \epsilon),\ -\mathbf{E}^{\pm}_z(\mathbf{x}, z_0 + \epsilon)\big),$$

where $\epsilon \in \mathbb{R}$ denotes the arbitrary shift from z_0 plane.

From definition (6) we for positive field have

$$\mathbf{E}^{+}(\mathbf{x}, z_0 + \epsilon) = \frac{1}{S_D} \sum_{n=-\infty}^{\infty} \int \frac{\big(\mathbf{x} - \mathbf{x}',\ (n_0 - 2n)L + \epsilon\big)}{\big(\|\mathbf{x} - \mathbf{x}'\|^2 + ((n_0 - 2n)L + \epsilon)^2\big)^{(D+1)/2}} \mathbb{P}(\mathbf{x}')\, \mathrm{d}\mathbf{x}'.$$

For the argument $z_0 - \epsilon$:

$$\mathbf{E}^{+}(\mathbf{x}, z_0 - \epsilon) = \frac{1}{S_D} \sum_{n=-\infty}^{\infty} \int \frac{\big(\mathbf{x} - \mathbf{x}',\ (n_0 - 2n)L - \epsilon\big)}{\big(\|\mathbf{x} - \mathbf{x}'\|^2 + ((n_0 - 2n)L - \epsilon)^2\big)^{(D+1)/2}} \mathbb{P}(\mathbf{x}')\, \mathrm{d}\mathbf{x}'.$$

In the last expression, change the summation index to $m = n_0 - n$. Then $n = n_0 - m$, and as n runs from $-\infty$ to ∞, so m also runs over all integers. We compute:

$$(n_0 - 2n)L - \epsilon = \big(n_0 - 2(n_0 - m)\big)L - \epsilon = (-n_0 + 2m)L - \epsilon = -\big((n_0 - 2m)L + \epsilon\big).$$

The denominator depends on the square of this expression and therefore remains unchanged. Consequently,

$$\mathbf{E}^{+}(\mathbf{x}, z_0 - \epsilon) = \frac{1}{S_D} \sum_{m=-\infty}^{\infty} \int \frac{\big(\mathbf{x} - \mathbf{x}',\ -\big((n_0 - 2m)L + \epsilon\big)\big)}{\big(\|\mathbf{x} - \mathbf{x}'\|^2 + ((n_0 - 2m)L + \epsilon)^2\big)^{(D+1)/2}} \mathbb{P}(\mathbf{x}')\, \mathrm{d}\mathbf{x}'.$$

Comparing with the expression for $+\epsilon$, we see that the x-components coincide completely, while the z-component changes sign. Case of negative field $\mathbf{E}^{-}(\widetilde{\mathbf{x}})$ can be considered similarly.

Note that in the particular case $\epsilon = 0$ the z-component of the field vanishes:

$$\mathbf{E}^{\pm}(\mathbf{x}, n_0 L) = \big(\mathbf{E}^{\pm}_x(\mathbf{x}, n_0 L),\ 0\big).$$

Now consider a field line that starts at a point $\widetilde{\mathbf{x}} = (\mathbf{x}, z)$ with $\mathbf{x} \in \operatorname{supp} \mathbb{P}$, $z \to 0^+$. Suppose the opposite of the theorem's statement holds, i.e. this field line leaves the region $z \in [0, L]$. Then, by continuity of electric field lines, there exists a point $\widetilde{\mathbf{x}}_0 = (\mathbf{x}_0, z_0)$, with $z_0 = 0$ or L, where this line crosses the boundary of the region. But at the point $\widetilde{\mathbf{x}}_0$ the z-component of the field is non-zero, which contradicts the symmetry property established above. This completes the proof of the first statement of the theorem.

We now prove the second statement. The electric field obeys Gauss's theorem; hence at the boundary (see, e.g., (Kolesov et al., 2025, Appendix B))

$$\mathbf{E}_z(\mathbf{x}, 0^+) - \mathbf{E}_z(\mathbf{x}, 0^-) = \mathbb{P}(\mathbf{x}).$$

By the reflection symmetry of the periodic configuration, $\mathbf{E}_z(\mathbf{x}, 0^-) = -\mathbf{E}_z(\mathbf{x}, 0^+)$, which yields

$$\mathbf{E}_z(\mathbf{x}, 0^+) = \frac{1}{2}\mathbb{P}(\mathbf{x}) \geq 0.$$

Analogously, at the negative plate,

$$\mathbf{E}_z(\mathbf{x}, L^-) = \frac{1}{2}\mathbb{Q}(\mathbf{x}) \geq 0.$$

These inequalities are strict on the supports of the respective distributions; outside the supports the corresponding one-sided boundary of z-component vanishes.

Now, note that the region $0 < z < L$ between the planes contains no charges; therefore the electric potential φ, related to the field by $\mathbf{E} = -\nabla\varphi$, satisfies Poisson's equation $\Delta\varphi = 0$. Each component of the field also satisfies the same equation:

$$\Delta \mathbf{E}_i = \Delta \left(-\frac{\partial}{\partial x_i}\right) \varphi = \left(-\frac{\partial}{\partial x_i}\right) \Delta\varphi = 0. \tag{12}$$

In particular, $\Delta \mathbf{E}_z = 0$, so $\mathbf{E}_z$ is a harmonic function. The compact support of the distributions guarantees that the field decays at infinity. Which allows one to apply the maximum and minimum principles for harmonic functions on the unbounded slab $z \in [0, L]$ (see, e.g., (Evans, 2010, Section 2.2)).

The weak maximum principle implies that a harmonic function attains its maximum and minimum on the boundary of the domain. Since $\mathbf{E}_z \geq 0$ on the boundaries $z = 0$ and $z = L$, it follows that $\mathbf{E}_z \geq 0$ everywhere inside the slab.

It remains to exclude the possibility that $\mathbf{E}_z = 0$ at some interior point. Suppose $\mathbf{E}_z(\widetilde{\mathbf{x}}) = 0$ for some $\widetilde{\mathbf{x}} = (\mathbf{x}, z)$ with $0 < z < L$. Because $\mathbf{E}_z \geq 0$, this is a global minimum. By the strong maximum principle, a harmonic function that attains its minimum in the interior must be constant throughout the domain. Hence $\mathbf{E}_z \equiv 0$ in the whole slab, which contradicts the strict positivity of $\mathbf{E}_z()$ on the supports of $\mathbb{P}$ and $\mathbb{Q}$ (via the continuity of the field up to the boundary). Therefore,

$$\mathbf{E}_z(\mathbf{x}, z) > 0 \qquad \text{for all } \mathbf{x} \in \mathbb{R}^D,\ 0 < z < L.$$

□

Theorem A.2 (C^{2}FM transport theorem). *Let $\mathbb{P}(\mathbf{x}^+)$ and $\mathbb{Q}(\mathbf{x}^-)$ be two data distributions with compact support, and let T be the deterministic map defined above. Let $\mathbf{x}^+$ be distributed as $\mathbb{P}(\mathbf{x}^+)$. Then $\mathbf{x}^- = T(\mathbf{x}^+)$ is distributed as $\mathbb{Q}(\mathbf{x}^-)$ almost surely:*

$$\textit{If } \mathbf{x}^+ \sim \mathbb{P}(\mathbf{x}^+) \Rightarrow T(\mathbf{x}^+) = \mathbf{x}^- \sim \mathbb{Q}(\mathbf{x}^-). \tag{13}$$

Proof. We verify that the field satisfies the interaction field requirements of (Manukhov et al., 2026, IFM). This automatically guarantees transport between the distributions.

- The first requirement of an interaction field is that a field line starting on a positive charge must almost surely (except, perhaps, a set of lines with zero flux) end on a negative charge. Because we work with electrostatic field lines, this requirement is fulfilled (see, e.g., (Kolesov et al., 2025, Appendix A)).
- The second requirement is flux conservation along a field-line tube between the planes $z \in (0, L)$. In this interval our charge configuration satisfies the requirement by Gauss's theorem (see the discussion in (Kolesov et al., 2025)).
- Finally, the generalized superposition principle reduces, as shown in (Manukhov et al., 2026), to the ordinary superposition principle (7) in the case of electrostatic fields.

Thus all IFM field requirements are met, which guarantees transport between the distributions. □

B. Technical Details

In this appendix, we describe the hyperparameters of our Algorithm 1 for different experiments in Table 1. Our code for the experiments is based on EFM's code `https://github.com/justkolesov/FieldMatching`.

We provide the implementation details for the synthetic experiments in §4. All tasks use two-dimensional Gaussian mixtures embedded into $\mathbb{R}^3$ with auxiliary coordinate z. The source distribution is placed at $z = 0$, the target distribution at $z = L$, and we set $L = 20$ in all experiments.

Two dimensional datasets are generated procedurally. The 5-GMM → 4-GMM task uses source means $(-8, 8), (8, 8), (0, 0), (-8, -8), (8, -8)$ and target means $(-8, 6), (8, 6), (-8, -6), (8, -6)$, with isotropic standard deviation 0.8 for all components. The 1-GMM → 2-GMM task uses source mean $(0, 0)$ with standard deviation 2.0, and target means $(-6, 0), (6, 0)$ with standard deviation 1.0. Mixture weights are uniform.

Experiment	D	Batch Size	L	C	NFE	Learning rate	W decay
1 GMM → 2 GMM §4.1	2	2000	20	2, 4, 20	400	2e-4	0.0
5 GMM → 4 GMM§4.1	2	2000	20	2,4,20	400	2e-4	0.0
MNIST digits 2→3 §4.2	3072	128	20	2, 4, 20	2000	2e-4	1e-4

Table 1. Hyper-parameters of Alg. 1 for the experiments, where D is the dimensionality of task, L is the distance between plates, NFE is number of inference steps, C is number of capacitors (see fig.1).

For the Image experiments (see §4.2), we adopt the Exponential Moving Average (EMA) approach with a decay rate of 0.99 to ensure smooth solutions, following common practices in prior work (Kolesov et al., 2025, EFM), (Xu et al., 2022; 2023, PFGM/PFGM++), (Ho et al., 2020, DDPM), and (Lipman et al., 2023, FM). We employ a linear learning rate scheduler that increases from 0 to 2×10^{-4} over the first $5{,}000$ iterations and then decreases monotonically. For optimization, we use the Adam optimizer with a learning rate of 2×10^{-4} and a weight decay of 1×10^{-4}.

For the image experiments (see §4.2), training our model takes less than 10 hours on a single NVIDIA A100 GPU with 30 GB of VRAM. We also report the peak GPU memory usage, which is approximately 10 GB during inference with a batch size of 128.

To evaluate inference time, we measure the time required to generate varying numbers of images (i.e., different batch sizes) for the MNIST 32×32 colored digits dataset (see §4.2) during the inference stage. The results are presented in Table 2. Notably, to ensure a fair comparison, we use the same Euler-based ODE solver, the same number of function evaluations (NFE), and the same neural architecture as the EFM method. Therefore, the inference time is identical across these methods.

	Batch Size = 256	Batch Size = 128	Batch Size = 64	Batch Size = 16	Batch Size = 1
time, sec	24.12	10.85	6.33	1.27	0.76

Table 2. The inference time with different $|B|$ batch sizes for generating colored digits from MNIST dataset 32x32 (see §4.2).