

Is There a Better Source Distribution than Gaussian?

Exploring Source Distributions for Image Flow Matching

Anonymous authors

Paper under double-blind review

Abstract

Flow matching has emerged as a powerful generative modeling approach with flexible source distribution choices. While Gaussian distributions are commonly used, the potential for better alternatives in high-dimensional data generation remains largely unexplored. In this paper, we propose a novel 2D simulation that captures high-dimensional geometric properties under the interpretable 2D setting, enabling us to analyze the learning dynamics of flow matching during training. Based on this analysis, we derive several key insights about flow matching behavior: (1) density approximation paradoxically degrades performance due to mode discrepancy, (2) directional alignment suffers from path entanglement when overly concentrated, (3) Gaussian’s omnidirectional coverage ensures robust learning, and (4) norm misalignment incurs substantial learning costs. Building on these insights, we propose a practical framework that combines norm-aligned training with directionally-pruned sampling. This approach maintains robust omnidirectional supervision essential for stable flow learning, while eliminating data sparse-region initializations during inference. Importantly, our pruning strategy can be applied to any flow matching model trained with a Gaussian source, providing immediate performance gains without the need for retraining. Empirical evaluations demonstrate consistent improvements in both generation quality and sampling efficiency. Our findings provide practical insight and guidelines for source distribution design and introduce a readily applicable technique for improving existing flow matching models.

1 Introduction

Flow Matching (FM) (Lipman et al., 2023; Liu et al., 2023b; Albergo et al., 2023) is a recently introduced approach for generative modeling that bridges ideas from Continuous Normalizing Flows (CNF) (Chen et al., 2018; Grathwohl et al., 2018) and diffusion models (Sohl-Dickstein et al., 2015; Song & Ermon, 2019; Ho et al., 2020; Nichol & Dhariwal, 2021; Song et al., 2021b). The core idea involves training a vector field so that trajectories, following an ordinary differential equation (ODE), transport samples from a simple base distribution to a complex target distribution, offering advantages in stable training, fast inference, and explicit likelihood estimation. Unlike diffusion models, which typically involve a fixed stochastic process that gradually adds Gaussian noise to data and requires learning a reverse denoising trajectory, flow matching offers complete freedom in defining the interpolation path between distributions. This flexibility not only enables the design of more efficient and straighter trajectories but also removes any constraint on the choice of source distribution.

This flexibility has motivated extensive research in enhancing FM performance through improved source-target pairing strategies (Tong et al., 2024; Pooladian et al., 2023; Davtyan et al., 2025) and developing straighter probability paths (Xing et al., 2023; Lee et al., 2023; Liu et al., 2023b) to reduce trajectory curvature and minimize function evaluations. Recent advances have leveraged prior knowledge of target distributions to construct better-aligned source distributions (Stark et al., 2023; Jing et al., 2024; Kolloviev et al., 2025), demonstrating substantial improvements through partial target information incorporation. For multimodal generation, recent works have exploited this source flexibility by using text embeddings directly as source distributions (Liu et al., 2024; He et al., 2025).

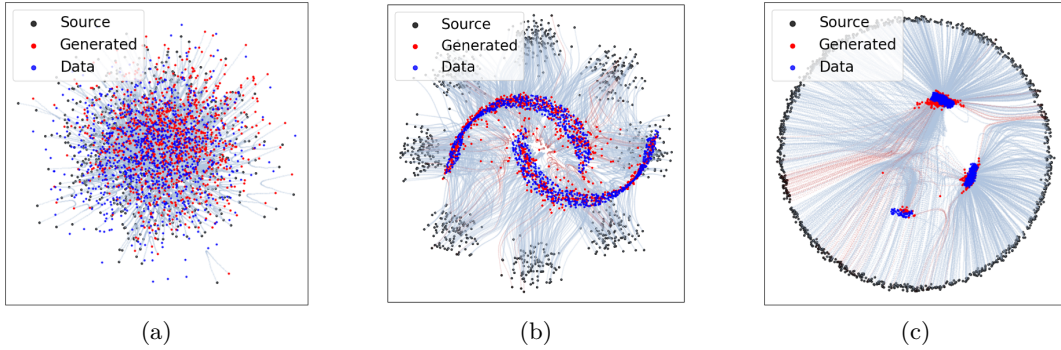


Figure 1: **2D simulations of flow matching.** The first two are examples of naive low-dimensional flow matching illustrations of (a) Gaussian to Gaussian and (b) 8 Gaussians to moons. Our proposed 2D simulation for flow matching is demonstrated in (c). Dots represent source (black), target data (blue), and generated samples (red). Blue dashed lines show successful ODE trajectories from source to generated samples, while red dashed lines indicate failed trajectories.

However, most existing target-aware source designs are still limited to relatively simple or low-dimensional datasets, such as time series (Kollovich et al., 2025) or protein-ligand interactions (Stark et al., 2023). To the best of our knowledge, there has been little to no prior work applying them to complex, high-dimensional data like natural images—likely due to the challenging nature of such applications. Moreover, naively replacing the standard Gaussian source with distributions from other modalities (*e.g.*, text embeddings) often fails to achieve stable target mappings, requiring extensive engineering as shown by prior work (Liu et al., 2024; He et al., 2025). This raises a fundamental question: What makes deviating the source from the Gaussian distribution so difficult, and what properties are required for a source to be advantageous in flow matching? To answer this, we begin with an intuitive hypothesis: source distributions that more closely approximate the density or geometric characteristics of the target distribution would lead to improved performance and faster convergence.

Since learning behaviors in high-dimensional datasets are difficult to interpret and visualize, we begin with low-dimensional experiment that enable clear visual analysis of the underlying learning dynamics. Several prior works have used naive low-dimensional experiments to demonstrate their effectiveness, *e.g.*, 2D Gaussian to 2D Gaussian in Fig. 1(a) or 8 Gaussians to moons in Fig. 1(b). However, we discover that merely using a low-dimensionality does not reflect the intricate but essential learning dynamics that emerge only in high-dimensional cases, and thus lessons learned from such a naive low-dimensional toy experiment do not fully transfer to the actual applications we are interested in, usually treating high-dimensional data. To address this, we redesign the low-dimensional experiments in Section 3 to better reflect the geometric properties by decomposing high-dimensional data through direction and norm components, as illustrated in Fig. 1(c). This 2D simulation allows us to more realistically analyze the learning dynamics of flow matching, preserving the essential characteristics of high-dimensional applications while enabling clear visualization of vector field behaviors in 2D.

Using this novel experiment framework, we examine in Section 4 how different source distribution strategies interact with various pairing methods (independent *vs.* in-batch optimal transport). This allows us to test our intuitive hypothesis that source distributions more closely approximating the density or geometric characteristics of the target distribution should improve performance. Specifically, we propose three ideas to redesign the source distribution: progressively approximating the density of the target distribution (*density approximation*), designing sources based on directional information of the target distribution (*directional alignment*), and aligning the expected norms of the source and target distributions (*norm alignment*).

From our simulation results and analysis in Sections 4 and 5, we derive several critical insights that reshape our understanding of source distribution design in flow matching: (1) Source distributions closely approximating the target density can paradoxically degrade performance, with stronger approximations leading to worse outcomes—a phenomenon we attribute to mode discrepancy where approximate sources omit low-density data

modes. (2) Directional alignment strategies, while theoretically appealing, often suffer from path entanglement due to insufficient source support and imperfect pairing, both of which hinder optimal performance. (3) The standard Gaussian source with independent pairing demonstrates unexpected robustness through omnidirectional coverage, while in-batch optimal transport pairing, despite its efficiency, sacrifices this crucial omnidirectional learning. (4) Regions with sparse or missing data consistently lead to generation failures due to insufficient supervision of the vector field during training. (5) A significant discrepancy between the norm distributions of the source and target incurs a substantial learning cost to resolve. These patterns hold consistently across both our controlled 2D simulations and real-world high-dimensional image datasets, demonstrating that our framework captures fundamental dynamics of flow matching.

Our key findings suggest that an effective source distribution design should leverage specific advantages while mitigating known weaknesses. Accordingly, we propose a new hybrid framework that combines training on a Gaussian source with “Norm Alignment” and employing “Pruned Sampling” during inference. This approach preserves the benefits of robust, omnidirectional learning from the Gaussian source (Finding 3) while resolving the learning inefficiency caused by norm mismatch (Finding 5). Furthermore, Pruned Sampling directly tackles the issue of passing through regions where the vector field is inaccurately learned (Finding 4) by excluding sampling from data-sparse regions, thus improving the quality of the final output. Notably, Pruned Sampling can be applied post-hoc to improve the performance of existing, pre-trained models without any need for retraining. This provides a practical pathway to immediately upgrade a wide range of flow matching models.

To summarize, we first propose a novel set of 2D simulation experiments designed to reveal and interpret the complex learning dynamics of flow matching in high-dimensional settings. Through these experiments, we derive critical findings that lead to a set of practical guidelines for source distribution design. Finally, we propose a readily applicable pruned sampling technique that can significantly improve existing flow matching models without the need for retraining. We believe these contributions not only advance the understanding of the flow matching field but also suggest new directions for developing more efficient and robust generative models.

2 Background and Related work

2.1 Flow Matching

Continuous normalizing flows (CNFs) define an ODE $dx = u_t(x)dt$ with a time-varying velocity field $u : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ that transports an initial distribution q_0 to a target distribution q_1 over $t \in [0, 1]$. This ODE induces a path of densities $p_t(x)$ satisfying the continuity equation $\partial_t p_t(x) + \nabla \cdot (p_t(x) u_t(x)) = 0$. Flow Matching (FM) (Albergo et al., 2023; Lipman et al., 2023; Liu et al., 2023b) is a simulation-free approach for training CNFs that avoids trajectory simulation by regressing the model’s vector field to a known probability flow. Instead of maximizing log-likelihood, FM assumes a prescribed probability path $p_t(x)$ that smoothly interpolates from $p_0 = q_0$ to $p_1 = q_1$, along with its corresponding velocity field $u_t(x)$. The neural ODE $v_\theta(x, t)$ is trained to match $u_t(x)$ via

$$\mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}_{t \sim \mathcal{U}[0,1], x \sim p_t(x)} \left[\|v_\theta(x, t) - u_t(x)\|^2 \right]. \quad (1)$$

FM is compatible with any interpolation path (e.g., linear or Gaussian) and enables scaling CNFs to high-dimensional data while maintaining stable optimization. However, this objective is often intractable due to the need to evaluate the marginal $p_t(x)$ and vector field $u_t(x)$. Conditional Flow Matching (CFM) addresses this by introducing a latent condition z to index a family of probability paths. Instead of a single fixed p_t , CFM considers a conditional density $p_t(x | z)$ with associated vector field $u_t(x|z)$. The overall path is recovered as $p_t(x) = \mathbb{E}_{z \sim q(z)} [p_t(x|z)]$. The CFM objective is given by

$$\mathcal{L}_{\text{CFM}}(\theta) = \mathbb{E}_{t \sim \mathcal{U}[0,1], z \sim q(z), x \sim p_t(x|z)} \left[\|v_\theta(x, t) - u_t(x|z)\|^2 \right]. \quad (2)$$

As it satisfies $\nabla_\theta \mathcal{L}_{\text{CFM}}(\theta) = \nabla_\theta \mathcal{L}_{\text{FM}}(\theta)$ under mild conditions, optimizing \mathcal{L}_{CFM} learns the correct global flow while sampling only from simpler conditional distributions. The key practical insight of CFM is to define

the latent variable z using the data samples themselves. In the standard independent coupling scheme, one sets $z = (x_0, x_1)$, where $x_0 \sim p_0(x)$ is a sample from the source distribution and $x_1 \sim p_1(x)$ is a sample from the target. The conditional path $p_t(x | x_0, x_1)$ thus becomes a simple “bridge” between a single source point and a single target point, making its path and vector field trivial to compute. Consequently, the training process only requires the ability to sample from $p_0(x)$ and $p_1(x)$, removing any need to evaluate its probability density (or log-density). This provides immense flexibility in the choice of p_0 .

With this simple framework, FM is able to generate more flexible probability paths (Chen & Lipman, 2024; Gat et al., 2024; Stark et al., 2024; Cheng et al., 2025; Kapusniak et al., 2024) than diffusion models (Song et al., 2021a; Song & Ermon, 2019; Song et al., 2021b; Rombach et al., 2022; Dhariwal & Nichol, 2021). FM also has been extended to various domains, including image (Esser et al., 2024; Dao et al., 2023; Ren et al., 2024), audio (Guan et al., 2024; Liu et al., 2023a; Prajwal et al., 2024), video (Jin et al., 2025; Polyak et al., 2024), molecule (Dunn & Koes, 2024; Song et al., 2023), and text generation (Hu et al., 2024).

2.2 Optimal Transport CFM (OT-CFM)

When trained with a globally optimal coupling, the flow map learned by FM aligns with the \mathcal{W}_2 geodesic between the source and target distributions. However, in practice, CFM typically adopts a simpler independent coupling $q_{\text{ind}}(x_0, x_1) = p_0(x_0)p_1(x_1)$, which pairs each source sample $x_0 \sim p_0$ with a randomly chosen target sample $x_1 \sim p_1$ (I-CFM). Given any coupling $q \in \Pi(p_0, p_1)$, where $\Pi(p_0, p_1)$ denotes the set of all valid joint distributions with marginals p_0 and p_1 , the transport cost and the optimal transport plan π are defined as:

$$\mathcal{C}(q) = \mathbb{E}_{(x_0, x_1) \sim q} [\|x_1 - x_0\|^2], \quad \pi = \arg \min_{q \in \Pi(p_0, p_1)} \mathcal{C}(q). \quad (3)$$

The independent coupling $q_{\text{ind}} = p_0(x_0)p_1(x_1)$ generally incurs a suboptimal transport cost. In this paper, we define the excessive cost relative to the optimal transport plan π as the *Wasserstein coupling gap*:

$$\Delta_W = \mathcal{C}(q_{\text{ind}}) - \mathcal{C}(\pi) \geq 0. \quad (4)$$

This gap quantifies the degree of misalignment between the randomly paired paths and the true geodesic paths. A large gap forces the learned flow to exhibit unnecessary curvature, deviating from the straightest possible transport map, which can in turn increase gradient variance and hinder training efficiency. To mitigate this issue, mini-batch optimal transport (BatchOT or OT-CFM) (Pooladian et al., 2023; Tong et al., 2024) selects, at each training step, the permutation that best matches the B source samples to the B target samples, yielding a locally optimal Wasserstein coupling. Because this permutation is recomputed for each mini-batch, the alignment does not persist across iterations, which can limit long-term convergence benefits.

2.3 Alternative Source Distributions

Target-Approximating Source Distribution. TSFlow (Kollovich et al., 2025) extends target-approximating source distributions to time series forecasting, integrating Gaussian Process priors within CFM to align source distribution with temporal target data structure. HarmonicFlow (Stark et al., 2023) employs physics-based priors for protein-ligand docking through 3D geometric constraints, utilizing protein backbone torsion-angle distributions to preserve SE(3)-equivariant properties. AlphaFlow (Jing et al., 2024) learns directly from PDB structural databases, replacing Gaussian noise with experimentally observed distributions to model biomolecular conformational variability. While promising, these approaches remain limited to simple domain-specific distributions and low-dimensional spaces.

Non-Gaussian Source in Cross-Modal Generation. CrossFlow (Liu et al., 2024) enables direct mapping from text to image embeddings using discrete language spaces as source distributions, overcoming traditional Gaussian space constraints. FlowTok (He et al., 2025) constructs a shared text-image token space by transforming Vision Transformer patch tokens into quantized discrete spaces, redefining image generation as 1D token sequence generation. However, these approaches require substantial engineering for source-target alignment, suggesting naive replacements may create more problems than solutions in complex high-dimensional data.

2.4 Von Mises-Fisher (vMF) Distribution

The von Mises-Fisher (vMF) distribution (Fisher, 1953), denoted as $\text{vMF}(\mu, \kappa)$, provides a mathematically principled framework for this directional concentration, defining a probability distribution on the unit hypersphere with mean direction μ and concentration parameter κ . As κ increases, the distribution becomes more sharply concentrated around the mean direction μ . When $\kappa = 0$, the distribution reduces to a uniform distribution over the unit sphere. This makes the vMF distribution particularly suitable for modeling directional data or sampling unit vectors with controlled angular concentration, which is useful in flow matching tasks where aligning the directionality of source and target samples is crucial.

3 Revisiting Gaussian Source Distributions for Flow Matching

The choice of the source distribution p_0 is a fundamental element in flow matching (Lipman et al., 2023; Liu et al., 2023b; Tong et al., 2024). While the standard isotropic Gaussian distribution $\mathcal{N}(0, I)$ is widely adopted for its simplicity and favorable mathematical properties, its optimality for complex image generation tasks remains as an open question. In this section, we point out geometric properties of Gaussian distributions, propose a functionally equivalent directional sampling paradigm, and derive key insights for designing effective source distributions with 2D simulations.

3.1 High-Dimensional Geometry and Directional Decomposition

The standard Gaussian $x_G \sim \mathcal{N}(0, I) \in \mathbb{R}^d$ serves as a conventional choice for the source distribution in generative models. In high-dimensional spaces, vast majority of independent and identically distributed (i.i.d.) Gaussian samples reside on a thin *hyperspherical shell*. Specifically, the Euclidean norm of x_G follows χ -distribution: $\|x_G\|_2 \sim \chi(d)$, whose mean and variance are approximately $\sqrt{d-1/2}$ and $1/2$, respectively. This indicates that for a sufficiently large d , most samples lie on a hypersphere whose radius is sharply concentrated in narrow band around $\sqrt{d-1/2}$. Motivated by this property, we propose the χ -Sphere decomposition, a directional decomposition scheme defined as:

$$x_G \approx r s_G, \quad \text{with } r \sim \chi(d), \quad s_G \sim \mathcal{U}(S^{d-1}), \quad (5)$$

where $\mathcal{U}(S^{n-1})$ denotes the uniform distribution on the unit sphere. Our χ -Sphere decomposition preserves all Gaussian statistics while explicitly factorizing each sample into an independent radius and a directional unit vector (see the derivation in Appendix A). To verify that this theoretical equivalence holds in practice, we further measure the FID scores of generated images using different combinations of training and inference source distributions. The result in Table 1 indicates that the samples from χ -Sphere and standard Gaussian distributions are interchangeable during both training and inference, without showing significant performance difference. Consequently, we regard these two sampling methods as functionally equivalent throughout this paper.

In practice, normalized data—typically shifted to zero mean and scaled to lie within $[-1, 1]$ —tend to concentrate within a bounded norm region, forming a thick *hyperspherical shell*. This enables a similar decomposition of data point $x_1 \equiv r_1 s_1 \in \mathbb{R}^d$, where $s_1 \in \mathbb{R}^d$ is the unit vector representing its direction and $r_1 \in \mathbb{R}$ is the norm of the data sample. Under this formulation, we can reasonably measure the “angular similarity” between a sample from Gaussian and a data point using the cosine similarity, $s_G \cdot s_1$. Viewing the data distribution through this directional perspective reveals it as a subset of directions of a full hypersphere, a viewpoint we explore further in Section 4.2.

Table 1: FID scores under different combinations of training and inference source distributions.

Training \ Inference	Gaussian	χ -Sphere
Gaussian	4.40	4.29
χ -Sphere	4.45	4.42

3.2 Proposed Simulation for High-Dimensional Flow Analysis

While previous works have visualized flow matching between simple 2D distributions, there has been little attention to understanding the learning dynamics of flow models in high-dimensional settings. [Davtyan et al. \(2025\)](#), for example, illustrates 2D Gaussian to 2D Gaussian, shown in Fig. [1\(a\)](#). [Tong et al. \(2024\)](#), illustrated in Fig. [1\(b\)](#), visualizes the transport paths from the 8 Gaussians to moons. These settings, however, fail to reflect the geometric properties of source and target samples and do not capture the behavior of flow matching in high-dimensional settings. To address this gap and analyze the behavior of flow matching models in high-dimensional settings, we propose a novel flow simulation setup that is more suitable for capturing the geometric properties uniquely observed in high-dimensional distributions.

Specifically, we simulate high-dimensional geometric structure in 2D by sampling random directions $\theta \in [0, 2\pi]$ and scaling with norms from $\chi(d)$ distributions (\bullet Source) following the χ -Sphere decomposition, as illustrated in Fig. [1\(c\)](#). When analyzing flow matching with the proposed density-approximated sources and directional sources introduced in Section [1](#), we normalize their mean norms to match the χ -Sphere radius. This normalization allows us to observe transport path patterns comparable to the χ -Sphere while ensuring similar geometric properties across all source types. For the target data (\bullet Data), we sample points along three clusters with varying densities, reflecting the common scenario where real high-dimensional datasets scaled to $[-1, 1]$ (e.g., CIFAR-10 with norm 27.2) typically have smaller norms than their corresponding Gaussian sources (55.4) and reside inside the concentration shell. We train flow matching models on these geometrically designed distributions and visualize their transport trajectories. Blue dashed lines (--- ODE trajectory) show successful paths from source (\bullet Source) to generated samples (\bullet Generated), while red dashed lines (--- Bad ODE trajectory) indicate failures. Here, we define success/failure based on whether the L2 distance from generated samples to the nearest data point is within one unit.

To complement our primary visual analysis, we report the Normalized Wasserstein [Balaji et al. \(2019\)](#), averaged over 10 runs, as an auxiliary measure. This metric captures how closely the generated samples resemble the true data distribution, even when the data contains multiple modes with imbalanced proportions. Unlike Wasserstein distance, the Normalized Wasserstein is robust to differences in mode proportions and focuses on the structural similarity between distributions. Lower values indicate better alignment with the true distribution. Since our setup is a controlled simulation, this metric should not be interpreted in absolute terms. Instead, it is most meaningful when used for relative comparison within the same group of source distributions. Additional metrics and implementation details are provided in Appendix [B](#).

4 Understanding Learning Dynamics of FM: Insights from 2D Simulations

Using the simulation approach proposed in Section [3.2](#), we conduct comprehensive experiments to verify the source distribution strategies proposed in Section [1](#). Although simplified, these experiments preserve the key geometric structures and transport dynamics of high-dimensional settings. They allow us to empirically probe the core mechanisms behind flow matching, offering insights into what drives its success or failure across different conditions.

4.1 Density Approximation Strategy

Recent studies ([Pooladian et al., 2023](#); [Tong et al., 2024](#)) have shown that improving the pairing between source and target distributions produce straighter probability flow trajectories, fewer sampling steps at inference, and expedited training. These improvements are largely attributed to reduced transport costs through better alignment. Motivated by these findings, we hypothesize that shaping the source distribution to closely resemble the target distribution—while preserving the improved pairing—might further reduce transport cost and improve generation performance.

To test our hypothesis, we first train a flow matching model, which we refer to as the *density approximator*, to transport samples from the χ -Sphere Gaussian distribution toward the target data distribution. The density approximator is trained for different number of iterations (specifically, 200, 6k, and 10k) to achieve progressively stronger approximations of the target distribution, and at each stage, we use the samples




Figure 2: **Visualization of flow matching with density-approximated source.** (a) OT-CFM with Approximated Source (iter 200), (b) OT-CFM with Approximated Source (iter 6000), (c) OT-CFM with Approximated Source (iter 10000). “Norm. W” denotes normalized Wasserstein, where a lower value indicates better generation performance.

generated by the partially trained model as the new approximated source distribution, denoted by \tilde{p}_0 . This approach allows us to observe how the progressively approximated source distribution \tilde{p}_0 acts as a new source distribution for flow matching.

Contrary to our initial hypothesis that a source distribution resembling the target would improve generation performance, we observe in Fig. 2 that more samples fall outside the data modes as the source approximates the target data. Also, higher Normalized Wasserstein is observed when the source distribution becomes increasingly similar to the target distribution. This indicates that the performance degrades in spite of a better density approximation between the source and the target, implying that a more similar source distribution might hinder accurate mode coverage and generation quality. This counterintuitive outcome challenges our hypothesis about source distribution design in flow matching and necessitates careful analysis of the underlying mechanism.

Mode Discrepancy. Taking a deeper look, as the source density approximator gradually learns to match the target distribution, the approximated source \tilde{p}_0 concentrate samples around each data mode (• Data). Subsequently, the distribution is adjusted to spread outward (• Source), aligning its norm with that of the χ -Sphere source, as discussed in Section 3.2. Comparing Fig. 2(a), (b) and (c), the density approximator reflects the two dominant data modes while gradually departing from the general spherical shell structure of the χ -Sphere. However, \tilde{p}_0 fails to capture the lower-left mode in sparse density regions even after 10,000 training iterations. This 2D simulation results reveal a fundamental problem that occurs in practice: any approximated distribution \tilde{p}_0 inevitably suffers from information loss, particularly in low-density regions where several modes may be underrepresented or omitted. This inaccurate approximation creates a pairing mismatch where certain target data points x_1 from underrepresented or omitted modes cannot be effectively paired with appropriate source samples x_0 from the approximated source \tilde{p}_0 . We term this phenomenon *mode discrepancy*. Consequently, even optimal transport pairing is forced to learn inefficient and complex trajectories, potentially increasing the learning difficulty instead of reducing it. This analysis reveals that mode discrepancy poses a fundamental limitation to density approximation strategies for flow matching source distributions.

4.2 Directional Alignment Strategy

Recognizing the critical importance of preserving all data modes during approximation, we shift our focus toward a different approach that sacrifices the norm information to better approximate data mode directions. Specifically, we leverage χ -Sphere decomposition and focus on the cosine similarity between the source unit vector s_G and the target unit vector s_1 . In an ideal scenario where we know all data mode directions, we can incorporate this knowledge into our source modeling. For each data point belonging to a specific mode

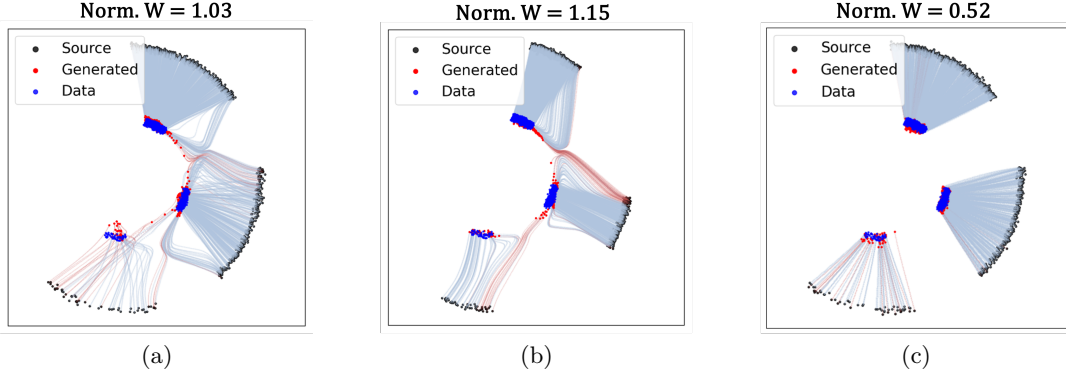


Figure 3: **Visualization of flow matching with ideally aligned directional source.** (a) OT-CFM with directional source, (b) OT-CFM with tight directional source, (c) global OT Pairing with directional source. ‘Norm. W’ denotes normalized Wasserstein.

direction, we construct a χ -Sphere source that maintains a cosine similarity above a threshold γ with the data (*i.e.*, $s_G \cdot s_1 \geq \gamma$). This approach allows us to avoid the mode discrepancy problem encountered earlier. This directional alignment approach is visualized under our experimental setting in Fig. 3, where the source successfully reflects all data mode directions. Combined with the improved pairing strategy, we expect this configuration to yield optimal results.

Imperfect Pairing. However, even when using a source distribution that covers all data modes through directional alignment and employing mini-batch OT pairing method, the configuration fails to achieve the expected optimal performance. As shown in Fig. 3(a) and (b), we observe that points are not generated perfectly within each data cluster—some red points scatter outside the blue target clusters. This suboptimal behavior contrasts with the global OT-paired directional alignment source in Fig. 3(c), where all generated points successfully converge to their respective data clusters. These findings highlight the limitation of mini-batch OT strategy that still falls significantly short of perfect global OT pairing.

Path Entanglement. In addition, interestingly, we observe that using a more tightly-aligned directional source by increasing γ in the directional source alignment actually leads to performance degradation, as shown in Fig. 3(b). From a geometric similarity perspective, as γ approaches 1, Fig. 3(b) shows that the source distribution becomes increasingly concentrated, spreading only as much as the data itself. However, we find that this reduction in source support adversely affects the overall performance. The underlying cause of this phenomenon can be accounted in Fig. 3(b), where paths from each source to target cluster become heavily entangled in localized regions during generation. This excessive entanglement leads to highly inconsistent vector field directions within local areas, making the vector field difficult to learn and unstable during optimization.

To better understand the cause of this entanglement, we provide a mathematical analysis of how the geometry of source-target pairings changes under increased concentration of the source distribution. Consider two target points x_1, x'_1 within the same mode cluster, with their associated source points x_0, x'_0 drawn from the corresponding directional source cluster. Their linearly interpolated paths are given by $x_t = (1-t)x_0 + tx_1$ and $x'_t = (1-t)x'_0 + tx'_1$, respectively, with separation: $d(t) = x_t - x'_t = (1-t)(x_0 - x'_0) + t(x_1 - x'_1)$. This separation obeys $\min_{t \in [0,1]} \|d(t)\| \leq \|x_0 - x'_0\| = \mathcal{O}((1-\gamma)^{1/2})$. Therefore, large values of γ cause the trajectories to be nearly coincident at initialization, which in turn leads to a sharp increase in the required local Lipschitz constant:

$$L_{\text{local}} \gtrsim \frac{\sin \theta}{\|x_1 - x'_1\| + \mathcal{O}((1-\gamma)^{1/2})}, \quad \text{where } \theta = \arccos \left(\frac{(x_1 - x_0)^\top (x'_1 - x'_0)}{\|x_1 - x_0\| \|x'_1 - x'_0\|} \right). \quad (6)$$

As a result, L_{local} becomes larger, making the optimization increasingly unstable and slow. More detailed derivation is provided in Appendix C.




Figure 4: **Visualization of flow matching with different methods.** (a) I-CFM with Gaussian source (b) OT-CFM with Gaussian source (c) I-CFM with pruned source.

These findings demonstrate that without oracle-level pairing, learning becomes significantly more difficult due to inconsistent vector fields induced by trajectory interference. To avoid this, the source distribution must retain the angular support sufficiently. The limitations observed here also extend to the density approximation discussed in Section 4.1, where increased approximation leads to reduced support, and imperfect pairing continues to pose challenges.

In conclusion, contrary to our initial expectation that combining mini-batch OT with improved source approximation (density or direction) would yield optimal results, our simulation reveal that increased approximation actually leads to diminished performance due to mode discrepancy, imperfect pairing mechanisms, and path entanglement. This counterintuitive finding underscores the delicate balance required between source distribution design and coupling strategies in flow matching approaches.

4.3 Pairing Method Analysis

To understand how different pairing methods affect the learning dynamics of flow models, we further analyze the generation trajectories and learning patterns of each approach.

Independent Pairing (I-CFM). The generation trajectory of I-CFM with χ -Sphere Gaussian source in Fig. 4(a) reveals patterns that can be interpreted as a conceptually two-stage transport process. In the initial stage, samples from the source move coarsely toward data-dense regions where the data modes are located. After this rough transport, fine-grained adjustments are performed in the second stage, where the paths are converged from all directions towards each mode. This occurs because the χ -Sphere source distribution provides uniform coverage across all directions. When paired independently with the target data, each data sample receives supervision from sources distributed omnidirectionally, enabling comprehensive vector field learning around each mode. This is clearly observed in Fig. 5(a), which shows the density of trajectories actually learned during training with independent pairing. The same principle applies to high-dimensional Gaussians, which also provide uniform directional coverage.

Improved Pairing (OT-CFM) In contrast, the OT-CFM trajectory heatmap in Fig. 5(b) shows a markedly different pattern. OT-CFM uses a batch-level optimal transport pairing strategy that minimizes the transport cost within each batch. This in-batch optimal pairing creates a more localized pairing where most source samples are matched with the nearby target cluster. Consequently, the model learns denser and more linear trajectories between optimally paired source and target samples. This can be seen in Fig. 4(b), where many source samples follow straighter paths compared to the more curved trajectories in Fig. 4(a) produced by I-CFM.

However, this improved path efficiency comes with a drawback. Because OT-CFM consistently pairs samples with nearby targets, the model learns the vector field primarily along narrow, cone-shaped directions centered around each data mode. As a result, it fails to learn the omnidirectional vector fields near the modes,

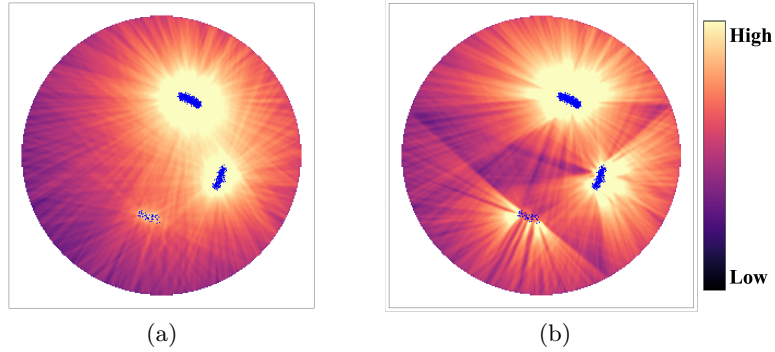


Figure 5: **Visualization of flow trajectory heatmap.** These heatmaps show the distribution of paths learned by each model. Color intensity corresponds to trajectory density, where brighter regions indicate higher concentrations of paths. (a) I-CFM and (b) OT-CFM.

which I-CFM captures better. This leads to insufficient learning of vector fields in broad angular directions, especially from the origin to the data modes, as shown in the trajectory heatmap in Fig. 5(b).

If all pairings were perfect, this focused learning approach would be both efficient and effective, as the model could learn the vector field accurately and consistently along these specific trajectories. In practice, however, the pairing is inevitably imperfect. When samples need to traverse these undertrained regions due to suboptimal pairings, they lack proper vector field guidance. Consequently, these samples often stall or veer off course, failing to reach the target data. This failure is clearly illustrated in Fig. 4(b), where multiple red points appear stranded midway, outside the data modes. This reveals a fundamental trade-off in OT-CFM: improved path efficiency comes at the expense of the robustness provided by omnidirectional coverage.

Low density Directions. A common observation in both Fig. 4(a) and (b) is that when samples are initialized in *low-density directions*—directions where data is absent and have large angular separation from data modes—they often fail to reach the data manifold during generation. These failed trajectories appear as clusters of red dashed lines (- -) near this direction. This issue arises because, interpolation paths from directions with sparse or absent data are sampled relatively less frequently during training, leading to lower training frequency for these regions. This is clearly visible in Fig. 5, where darker regions (such as the left and lower left areas) indicate lower training frequency compared to other parts of the heatmap. The rarity of path traversal during training results in insufficient supervision for the velocity field in these regions. Consequently, the vector field in such areas becomes inaccurately learned, increasing the likelihood of generation failures from samples initialized in low-density directions.

Building upon these insights, we experiment with a simple strategy that reflects the flow learning dynamics observed in our simulation. By comparing the two pairing methods, we found that learning an omnidirectional vector field around each data mode leads to a more robust velocity field. We also observed that, across all pairing strategies, source samples from low-density directions tend to produce poor generations. Combining these findings, we train I-CFM using the full χ -Sphere source to encourage robust and omnidirectional learning. Then, during sampling, we prune source samples from low-density directions to avoid poorly trained regions. This hybrid strategy, which is visualized in Fig. 4, resulted in significantly fewer failed trajectories and improved performance by maintaining comprehensive training coverage while guiding sampling toward more reliable paths during generation. Through this approach, we gained valuable insights that could help identify better source distributions for flow matching models.

5 Empirical Validation on High-Dimensional Image Datasets

In this section, we extend our analysis to high-dimensional image datasets through parallel experiments that validate the hypotheses and findings from Section 4. We investigate whether the insights from our low-dimensional simulations transfer to real-world image generation tasks.

Table 2: **Comparison of Approximated Target Distribution.** Unconditional generation results on CIFAR10 using different source distributions. All models are trained with OT-CFM (Tong et al., 2024).

	Gaussian	DCT	GMM-1	GMM-2	GMM-10	CNF
FID ↓	4.40	4.20	11.75	12.49	12.11	17.18

5.1 Density Approximation Strategy

In Section 4.1, we examined the intuitive hypothesis that using source distributions that better approximate the target distribution’s density would improve the efficiency and performance of flow matching through our 2D simulations. Our findings revealed that density-approximated sources introduce mode discrepancy issues that hinder effective training and yield inferior performance compared to baseline. To evaluate whether this conclusion in 2D simulation can be generalized to high-dimensional real-world datasets, we investigate three progressively expressive approximations to the target distribution and assess their impact on flow matching performance and training dynamics. To evaluate progressively stronger approximations to the target distribution, we employ the following density approximation methods:

- **Discrete Cosine Transform (DCT)** (Ahmed et al., 1974; Strang, 1999), commonly used in image compression (*e.g.*, JPEG (Wallace, 1992)), separates spatial frequencies of the image, discarding high frequency components that are less perceptible to humans (Wang et al., 2004). Inspired by this, we apply DCT filtering to Gaussian source to remove less prominent high-frequency components of the target distribution. See Appendix D for more details.
- **Gaussian Mixture Model (GMM)** (Reynolds et al., 2009) approximates a complex distribution by fitting multiple Gaussians. We train GMMs with 1, 2, and 10 components using the Expectation-Maximization (EM) algorithm (Moon, 1996), and use these as alternative source distributions.
- **Continuous Normalizing Flow (CNF)** (Grathwohl et al., 2018; Chen et al., 2018) constructs flexible distributions by continuously transforming a simple base density through a series of invertible mappings parameterized by neural ODEs. Although its invertibility constraint might limit expressiveness, CNF can still serve as a sophisticated approximation of the data distribution. We train FFJORD (Grathwohl et al., 2018), a popular CNF model, and use its output as the source distributions for flow matching.

As shown in Table 2, while DCT-based mild refinement improves performance compared to the Gaussian baseline, stronger approximations to the target distribution progressively degrade performance (DCT > GMM > CNF in terms of quality). This finding aligns with our simulation results from Section 4, where we observed similar performance degradation with stronger density approximations.

Although high-dimensional settings make it difficult to analyze the specific underlying mechanisms, the consistent trend across both simulations and real experiments suggests a fundamental issue with density-based approximation strategies. Drawing from our 2D analysis, a key insight is that approximating the target distribution inevitably leads to information loss, particularly in low-density regions. These regions tend to be underrepresented or omitted in the support of the approximate source \hat{p}_0 . As demonstrated in our 2D simulation (Fig. 2(b)), when no source sample $x_0 \sim \hat{p}_0$ is directionally aligned with a given target sample $x_1 \sim p_1$ (such as samples from sparse modes), even OT-based coupling produces inefficient, entangled trajectories that increase training complexity and hurt generative performance.

5.2 Directional Alignment Strategy

In our previous analysis, we identified a fundamental limitation: approximating the target distribution’s density inevitably leads to information loss, and using such approximations as source distributions results in mode discrepancy, where certain data modes lack corresponding suitable samples, hindering effective training. To address this challenge, we propose a directional alignment strategy that focuses on comprehensive mode coverage, paralleling our 2D simulations in Section 4.2.

Table 3: **Directional Alignment Performance.** FID scores on CIFAR10 using Euler integration for 100 NFE, comparing standard Gaussian source with directionally-aligned alternatives.

Algorithm	Oracle-vMF			Kmeans-vMF					Gaussian
	$\kappa = \infty$	$\kappa = 3000$	$\kappa = 1500$	$\kappa = 1500$	$\kappa = 900$	$\kappa = 300$	$\kappa = 100$	$\kappa = 50$	($\kappa = 0$)
OT-CFM	0.74	1.98	2.86	7.11	5.86	4.64	4.22	4.15	4.40

We find that the FID between normalized and original data is only 0.29, demonstrating that directional information is the key factor for generation quality. Based on this insight, rather than attempting to match the complete target density—which suffers from information loss in sparse regions—we focus on preserving directions where target data exists while removing directions without data. This approach prioritizes directional information of data modes while ignoring norm information, thereby avoiding the mode discrepancy problems inherent in density approximation strategies.

To obtain a distribution aligned with target directions, we leverage the χ -Sphere decomposition introduced in Section 3.1 and vMF distribution. This allows us to approximate directional pairing by generating source samples x_0 that are directionally aligned with each target sample x_1 . The specific sampling strategy follows:

$$x_0 = r s_0, \quad \text{with } r \sim \chi(d), \quad s_0 \sim \text{vMF}(\mu(x_1), \kappa(x_1)), \quad (7)$$

where the vMF’s mean direction $\mu(x_1)$ aligns the direction of x_0 with x_1 , while the probabilistic radius r and concentration parameter $\kappa(x_1)$ enable probabilistic rather than deterministic pairing, providing robust alignment through controlled variability.

5.2.1 Oracle Approach

To investigate the potential performance when sampling is focused on directions where data actually exists, we design an oracle pairing experiment. In this oracle scenario, the vMF mean direction is set to each normalized data point, and varying κ values control the tightness of pairing between source and target samples. This experimental setting closely corresponds to the ideal scenario illustrated in our 2D simulation (Fig. 3(c)). When κ approaches infinity, x_0 and x_1 become directionally identical, reducing the task to simple norm matching. Conversely, when $\kappa = 0$, the distribution becomes equivalent to Gaussian according to Section 3.1

As shown in Table 3, while κ approaching infinity yields FID scores near zero, this result is trivial since it merely reproduces the training data with norm scaling. More importantly, at reasonable κ values (1500, 3000), we observe that multiple generations from the same $\text{vMF}(u_1, \kappa)$ distribution produce images that are structurally similar yet vary in fine details. This demonstrates that the model learns to generate diverse samples around each data point rather than simply memorizing individual training examples (see Fig. 7 in appendix). Even under oracle conditions, the model generates novel images not present in the original distribution, confirming the potential of directional alignment strategies.

However, this approach has a critical limitation: it requires storing the entire training dataset to generate corresponding direction-aligned source samples, making it impractical for large-scale applications.

5.2.2 Clustering-based Approach

The oracle pairing method, while effective, is impractical due to its massive storage requirements. To address this, we propose a more scalable alternative: clustering-based pairing. This method approximates the oracle setup by grouping similar data points and assigning them a shared source distribution, drastically reducing the computational and storage overhead.

The process is straightforward. First, we partition the normalized target data into K groups using spherical K -means clustering (MacQueen, 1967). For each of the K clusters, we then construct a vMF distribution as $\text{vMF}(c_j / \|c_j\|, \kappa)$, where c_j is the centroid of cluster j , and sample source points following Eq. (7). Each target sample x_1 is then paired with a source sample x_0 drawn from the vMF corresponding to its assigned cluster. While this pairing is suboptimal compared to the oracle pairing, it offers a practical compromise between random Gaussian pairing and perfect directional alignment with proper K and κ .

In Table 3, we benchmark the performance across a range of concentrations κ , fixing $K = 3$ found by an elbow study with silhouette scores. As the concentration parameter κ decreases, each distribution covers a broader angular range. When $\kappa = 0$, the distribution becomes nearly identical to Gaussian, and generation performance is correspondingly same. For moderately concentrated ranges of $\kappa \in \{50, 100\}$, we observe notable performance improvements. At high concentrations ($\kappa \geq 300$), however, performance drops below that of the Gaussian baseline. This degradation, similar to what we observed in Fig. 3(b), stems from path entanglement due to support deficiency as discussed in Section 4.2. When the source distribution becomes too concentrated, trajectories from nearby source points become entangled, making the vector field difficult to learn and causing training instability.

This finding highlights a fundamental trade-off: while directional alignment offers advantages, the source distribution must retain sufficient angular support to avoid trajectory interference. Without oracle-level pairing that guarantees consistent and perfect alignment, overly concentrated distributions create more problems than they solve, confirming that robust flow learning requires balanced coverage rather than extreme concentration.

6 Proposed Method

In this section, we introduce two complementary strategies based on previous findings. We first propose “Pruned Sampling”, which creates a more efficient source distribution by identifying and removing directions that are directionally irrelevant to the target data. We then introduce “Norm Alignment” to resolve the scale mismatch between the source and target distributions. By creating a data-informed source, these methods improve the performance and stability of flow matching models without requiring architectural changes or retraining.

6.1 Pruned Sampling

Our previous experiments showed that we need to maintain sufficient support, but it is difficult to determine how much support is needed. Instead of attempting to find this amount, we propose the opposite strategy: pruning the Gaussian. The core idea is to start with a Gaussian source that covers all directions, then identify and prune away regions that are directionally irrelevant to the target data. By removing directions where data is absent, we can adjust the source distribution to better align with the target manifold.

While several approaches can be employed to identify directions where data is absent or sparse, we utilize Principal Component Analysis (PCA) (Pearson, 1901) for simplicity. Specifically, given a dataset $\mathcal{D} = \{x_i\}_{i=1}^N \subset \mathbb{R}^d$, we first L_2 -normalize each sample by $\tilde{x}_i = x_i / \|x_i\|_2$ and compute an orthonormal basis $V = [v_1, \dots, v_d] \in \mathbb{R}^{d \times d}$ via PCA, where the basis vectors are ordered by their corresponding eigenvalues in descending order. To account for directional symmetry, we construct an extended basis by including the opposite directions:

$$V_{\text{ext}} = [v_1, \dots, v_d, v_{d+1}, \dots, v_{2d}] \in \mathbb{R}^{d \times 2d}, \quad (8)$$

where $v_{k-d} = -v_k$ for $k = d+1, \dots, 2d$. For every basis vector $v_k \in V_{\text{ext}}$, we measure the largest cosine similarity $\gamma_k = \max_i \langle v_k, \tilde{x}_i \rangle$ for $k = 1, \dots, 2d$. Directions with a low γ_k indicate that they are remote from the data distribution. We denote a set of such basis vectors $\mathcal{R} = \{v \in V_{\text{ext}} : \gamma_t \leq \tau\}$, where τ is a hyperparameter indicating the threshold. We then apply rejection sampling; that is, we sample $x_0 \sim \mathcal{N}(0, I)$ and retain it and only if $\max_{v \in \mathcal{R}} \langle v, x_0 / \|x_0\|_2 \rangle \leq \tau_r$, where $\tau_r > \tau$ provides a more conservative margin to prevent potential support deficiency in critical directions. Detailed experimental settings can be found in Appendix E.

Based on this approach, we compare three strategic approaches: (i) Gaussian→Gaussian, where the model is trained and sampled from the complete Gaussian distribution; (ii) Pruned→Pruned, which involves training and sampling exclusively from directions that meet the pruning criteria; and (iii) Gaussian→Pruned where the model is trained on the complete Gaussian distribution but sampled only from its pruned subset.

Empirical results in Table 4 reveal unexpected patterns. The fully pruned method (ii) yields inconsistent performance compared to the Gaussian baseline (i). While it offers marginal improvements at a high number of function evaluations (NFE), it has minimal impact on I-CFM at low NFE and negatively affects the

Table 4: **Pruned Sampling.** Comparison of FID scores for ICFM and OTCFM under different training and inference source configurations on CIFAR10. Best scores are indicated in **bold**.

Method	Train Source	Inference Source	FID (NFE=5)	FID (NFE=100)
I-CFM	Gaussian	Gaussian	34.74	4.36
	<i>Pruned</i>	<i>Pruned</i>	34.71	4.17
	Gaussian	<i>Pruned</i>	28.99	3.95
OT-CFM	Gaussian	Gaussian	18.55	4.40
	<i>Pruned</i>	<i>Pruned</i>	20.47	4.18
	Gaussian	<i>Pruned</i>	17.29	4.10

Table 5: FID scores on ImageNet64 using Euler integration for OT-CFM w/wo source *Pruned*.

Model	Source	NFE			
		5	10	20	100
OT-CFM	Gaussian	53.95	18.84	10.62	9.10
	<i>Pruned</i>	49.56	16.70	9.54	8.78

performance of OT-CFM. In contrast, the hybrid method (iii) consistently outperforms the other methods across all evaluated configurations.

Although this observation may seem counterintuitive, it aligns with our 2D simulation findings in Section 4.3. The effectiveness of the hybrid “Gaussian→Pruned” strategy stems from the distinct requirements of the training and inference phases.

During training, leveraging the full Gaussian source is advantageous. A high-dimensional Gaussian distribution provides uniform angular coverage, ensuring the model learns the vector field omnidirectionally around data modes. This comprehensive training is critical for building a robust model that can generalize across diverse initialization conditions. Conversely, restricting training to only the pruned directions (\mathcal{S}) limits the model’s exposure. Without supervision from the excluded directions ($\mathbb{S}^{d-1} \setminus \mathcal{S}$), the model’s ability to generalize to unseen trajectories is compromised. This limitation is particularly detrimental at low NFEs, where the model has fewer opportunities to correct for poor initializations.

During inference, however, certain regions of the Gaussian source are suboptimal for sampling. These include areas devoid of data and singular points equidistant from multiple data modes. Source samples (x_0) from these regions are surrounded by fewer interpolation paths, leading to less frequent updates during training (see Fig. 5). Consequently, the vector field near these samples is learned less accurately, making them prone to generating erroneous outputs. Pruning these directions from source distribution during inference effectively guides the sampling process toward more reliable paths, enhancing generation quality.

In summary, the optimal strategy involves training on the complete Gaussian distribution to ensure the model learns a comprehensive and generalizable vector field, while sampling exclusively from the pruned subset to avoid regions where the learned field is likely to be inaccurate.

This finding has important practical implications. Pruned Sampling can be applied as a post-processing to *any* pretrained flow matching model that uses a Gaussian source distribution, requiring no retraining or architectural modifications. We demonstrate the effectiveness and scalability of our approach across multiple settings. In Table 5, applying Pruned Sampling to a higher-dimensional ImageNet64 pretrained model yields consistent performance improvements, confirming the scalability of our approach. Furthermore, Table 6 shows that Pruned Sampling consistently enhances performance for both OT-CFM and I-CFM, regardless of the number of function evaluations (NFE). By simply filtering out initialization points from data-sparse regions during inference, existing models can achieve immediate performance improvements with minimal computational overhead.

Table 6: FID scores on CIFAR10 using Euler integration for various numbers of function evaluations (NFE) for OT-CFM and I-CFM, comparing the baseline (Gaussian source) with our proposed method that incorporates *Pruned* and/or *NormAlign*. Our methods are indicated in *italics*.

Train Method	Source	NFE				
		5	10	20	40	100
OT-CFM (Baseline)	Gaussian	18.55	10.50	7.14	5.48	4.40
OT-CFM	<i>Pruned</i>	17.29	9.34	6.39	4.80	4.11
OT-CFM + <i>NormAlign</i>	Gaussian	24.66	11.33	6.91	5.01	4.15
OT-CFM + <i>NormAlign</i>	<i>Pruned</i>	21.96	9.81	5.95	4.57	3.87
I-CFM (Baseline)	Gaussian	34.74	13.22	7.67	5.55	4.36
I-CFM	<i>Pruned</i>	28.99	10.44	6.38	4.83	3.95
I-CFM + <i>NormAlign</i>	Gaussian	52.23	18.78	8.12	4.85	3.82
I-CFM + <i>NormAlign</i>	<i>Pruned</i>	48.44	16.82	7.10	4.56	3.64

6.2 Norm Alignment Strategy

In addition to the directional considerations, we identify another fundamental disparity in flow matching: the significant difference in scale between the source and target distributions. Flow matching models conventionally operate between a standard Gaussian source $\mathcal{N}(0, I)$ and the target data distributions, creating a substantial norm disparity that impacts model performance. The standard Gaussian source samples typically exhibit norms of approximately \sqrt{d} , while target data samples (often normalized to $[-1, 1]^d$) possess norms considerably smaller. This disparity necessitates significant computational resources—either in model capacity or training iterations—to resolve, potentially diverting attention from more crucial aspects of distribution estimation.

To address this disparity, we propose a *Norm Alignment* strategy. Specifically, we compute the expected norm of the target distribution, $E_{x_1 \sim p_1}[\|x_1\|]$, and that of the source distribution, $E_{x_0 \sim p_0}[\|x_0\|]$, which is $E[\chi(d)]$ when $p_0 = \mathcal{N}(0, I)$. We scale the target samples to $x'_1 \equiv x_1 \cdot E[\|x_0\|]/E[\|x_1\|]$, placing the transformed target distribution p'_1 on a hypersphere with a radius matching the average norm of p_0 . During inference, we reverse this scaling by multiplying x_1 by $E[\|x_1\|]/E[\|x_0\|]$, thus recovering the original scale of the target samples.

Aligning the average norms of source and target samples through the proportional scaling described above yields substantial performance gains, as demonstrated in Table 6. This suggests that flow models face greater difficulty than expected in learning to resolve norm mismatches, requiring significant computational resources that could otherwise be devoted to more essential aspects of the generation task. However, at low numbers of function evaluations (NFE), applying *Norm Alignment* can actually degrade performance. This is because, when both source and target distributions lie on the same norm hypersphere, some transport trajectories are restricted to move close to the surface of this hypersphere. As a result, the transport paths become more curved, requiring a greater number of NFEs for the model to accurately follow these paths and achieve effective transport. In contrast, without norm alignment, source samples can move more directly toward lower-norm regions near the data manifold, resulting in straighter transport paths at low NFEs.

As a final remark, combining *Norm Alignment* and *Pruned Sampling* at NFE 100 on CIFAR10 datasets yields substantial improvements: reducing FID by 0.67 for OT-CFM and 0.72 for I-CFM. These results highlight the importance of addressing both directionality and norm alignment in the source distribution to fully unlock the potential of flow matching models.

7 Conclusion

This work investigates whether alternative source distributions can outperform the standard Gaussian in flow matching. To better understand how flow matching models learn, we propose a novel 2D simulation designed to visualize and interpret high-dimensional behaviors in flow matching. Our simulations show strong

agreement with actual high-dimensional results, highlighting the effectiveness of this approach for analyzing and gaining insights into flow matching.

Our analysis reveals that intuitive strategies often fail. Density approximation approaches, using models like GMMs or CNFs, suffer from mode discrepancy; by omitting low-density target modes, they force the model to learn inefficient pathways. Similarly, directional alignment methods induce path entanglement when the source becomes too concentrated, destabilizing optimization. We find that the success of the Gaussian distribution lies in its omnidirectional coverage, which ensures robust vector field learning around all data modes.

Building on these insights, we introduce a hybrid framework that combines Norm Alignment during training with Pruned Sampling at inference. This preserves the robustness of Gaussian-based training while eliminating problematic initializations from data-sparse regions. Crucially, Pruned Sampling can be applied to pre-trained models without any retraining, offering a significant practical advantage. Our extensive experiments confirm consistent improvements across multiple settings, demonstrating both the practical value of our insights and the effectiveness of our readily applicable technique for advancing flow matching.

References

- N. Ahmed, T. Natarajan, and K.R. Rao. Discrete cosine transform. *IEEE Transactions on Computers*, C-23(1):90–93, 1974. doi: 10.1109/T-C.1974.223784.
- 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.
- Yogesh Balaji, Rama Chellappa, and Soheil Feizi. Normalized wasserstein for mixture distributions with applications in adversarial learning and domain adaptation. In *ICCV*, 2019.
- Ricky TQ Chen and Yaron Lipman. Flow matching on general geometries. In *ICLR*, 2024.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *NeurIPS*, 2018.
- Chaoran Cheng, Jiahua Li, Jiajun Fan, and Ge Liu. α -flow: A unified framework for continuous-state discrete flow matching models. *arXiv preprint arXiv:2504.10283*, 2025.
- Quan Dao, Hao Phung, Binh Nguyen, and Anh Tran. Flow matching in latent space. *arXiv preprint arXiv:2307.08698*, 2023.
- Aram Davtyan, Leello Tadesse Dadi, Volkan Cevher, and Paolo Favaro. Faster inference of flow-based generative models via improved data-noise coupling. In *ICLR*, 2025.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009. doi: 10.1109/CVPR.2009.5206848.
- Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- Ian Dunn and David Ryan Koes. Mixed continuous and categorical flow matching for 3d de novo molecule generation. *ArXiv*, pp. arXiv–2404, 2024.
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. URL <https://arxiv.org/abs/2403.03206>, 2, 2024.
- Ronald Aylmer Fisher. Dispersion on a sphere. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 217(1130):295–305, 1953.

- Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and Yaron Lipman. Discrete flow matching. In *NIPS*, 2024.
- Will Grathwohl, Ricky TQ Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. Ffjord: Free-form continuous dynamics for scalable reversible generative models. In *ICLR*, 2018.
- Wenhao Guan, Kaidi Wang, Wangjin Zhou, Yang Wang, Feng Deng, Hui Wang, Lin Li, Qingyang Hong, and Yong Qin. Lafma: A latent flow matching model for text-to-audio generation. *arXiv preprint arXiv:2406.08203*, 2024.
- Ju He, Qihang Yu, Qihao Liu, and Liang-Chieh Chen. Flowtok: Flowing seamlessly across text and image tokens. *arXiv preprint arXiv:2503.10772*, 2025.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 33:6840–6851, 2020.
- Vincent Hu, Di Wu, Yuki Asano, Pascal Mettes, Basura Fernando, Björn Ommer, and Cees Snoek. Flow matching for conditional text generation in a few sampling steps. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 380–392, 2024.
- Yang Jin, Zhicheng Sun, Ningyuan Li, Kun Xu, Hao Jiang, Nan Zhuang, Quzhe Huang, Yang Song, Yadong Mu, and Zhouchen Lin. Pyramidal flow matching for efficient video generative modeling. In *ICLR*, 2025.
- Bowen Jing, Bonnie Berger, and Tommi Jaakkola. Alphafold meets flow matching for generating protein ensembles. In *ICML*, 2024.
- Kacper Kapusniak, Peter Potaptchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael Bronstein, Joey Bose, and Francesco Di Giovanni. Metric flow matching for smooth interpolations on the data manifold. In *NIPS*, 2024.
- Marcel Kollovich, Marten Lienen, David Lüdke, Leo Schwinn, and Stephan Günnemann. Flow matching with gaussian process priors for probabilistic time series forecasting. In *ICLR*, 2025.
- Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical Report 0, University of Toronto, Toronto, Ontario, 2009. URL <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>.
- Sangyun Lee, Beomsu Kim, and Jong Chul Ye. Minimizing trajectory curvature of ode-based generative models. In *ICML*, 2023.
- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *ICLR*, 2023.
- Alexander H Liu, Matt Le, Apoorv Vyas, Bowen Shi, Andros Tjandra, and Wei-Ning Hsu. Generative pre-training for speech with flow matching. *arXiv preprint arXiv:2310.16338*, 2023a.
- Qihao Liu, Xi Yin, Alan Yuille, Andrew Brown, and Mannat Singh. Flowing from words to pixels: A framework for cross-modality evolution. *arXiv preprint arXiv:2412.15213*, 2024.
- Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *ICLR*, 2023b.
- James MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, volume 5, pp. 281–298. University of California press, 1967.
- Todd K Moon. The expectation-maximization algorithm. *IEEE Signal processing magazine*, 13(6):47–60, 1996.

- Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *ICML*, 2021.
- Karl Pearson. Principal components analysis. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 6(2):559, 1901.
- A Polyak, A Zohar, A Brown, A Tjandra, A Sinha, A Lee, A Vyas, B Shi, CY Ma, CY Chuang, et al. Movie gen: A cast of media foundation models, 2025. *arxiv:2410.13720*, 2024.
- Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman, and Ricky TQ Chen. Multisample flow matching: Straightening flows with minibatch couplings. In *ICML*, 2023.
- KR Prajwal, Bowen Shi, Matthew Lee, Apoorv Vyas, Andros Tjandra, Mahi Luthra, Baishan Guo, Huiyu Wang, Triantafyllos Afouras, David Kant, et al. Musicflow: Cascaded flow matching for text guided music generation. *arXiv preprint arXiv:2410.20478*, 2024.
- Sucheng Ren, Qihang Yu, Ju He, Xiaohui Shen, Alan Yuille, and Liang-Chieh Chen. Flowar: Scale-wise autoregressive image generation meets flow matching. *arXiv preprint arXiv:2412.15205*, 2024.
- Douglas A Reynolds et al. Gaussian mixture models. *Encyclopedia of biometrics*, 741(659-663):3, 2009.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pp. 234–241. Springer, 2015.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, pp. 2256–2265. pmlr, 2015.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021a.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *NeurIPS*, 32, 2019.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*, 2021b.
- Yuxuan Song, Jingjing Gong, Minkai Xu, Ziyao Cao, Yanyan Lan, Stefano Ermon, Hao Zhou, and Wei-Ying Ma. Equivariant flow matching with hybrid probability transport for 3d molecule generation. *Advances in Neural Information Processing Systems*, 36:549–568, 2023.
- Hannes Stark, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. Harmonic prior self-conditioned flow matching for multi-ligand docking and binding site design. In *NeurIPS 2023 AI for Science Workshop*, 2023.
- Hannes Stark, Bowen Jing, Chenyu Wang, Gabriele Corso, Bonnie Berger, Regina Barzilay, and Tommi Jaakkola. Dirichlet flow matching with applications to dna sequence design. In *ICML*, 2024.
- Gilbert Strang. The discrete cosine transform. *SIAM review*, 41(1):135–147, 1999.
- 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. *TMLR*, 2024.
- Gregory K Wallace. The jpeg still picture compression standard. *IEEE transactions on consumer electronics*, 38(1):xviii–xxxiv, 1992.

Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

Siyu Xing, Jie Cao, Huaibo Huang, Xiao-Yu Zhang, and Ran He. Exploring straighter trajectories of flow matching with diffusion guidance. *arXiv preprint arXiv:2311.16507*, 2023.