# Isokinetic Flow Matching for Pathwise Straightening of Generative Flows

**Tauhid Khan** [1]

## Abstract

Flow Matching (FM) constructs linear conditional probability paths, but the learned marginal velocity field exhibits curvature due to trajectory superposition, inflating numerical truncation errors and bottlenecking few-step sampling. We introduce **Isokinetic Flow Matching (Iso-FM)**, a lightweight, Jacobian-free regularizer that penalizes pathwise acceleration via a self-guided finite-difference approximation of the material derivative $Dv/Dt$. Operating as a plug-and-play addition to single-stage FM training, Iso-FM requires only standard forward evaluations and stop-gradient targets. On CIFAR-10 (DiT-S/2), Iso-FM reduces conditional non-OT FID@2 from 78.82 to 27.13, a 2.9× relative efficiency gain and achieves a best-observed FID@4 of 10.23. These results demonstrate that acceleration regularization is a principled, compute-efficient mechanism for improving the quality–NFE trade-off in flow-based generative models.

## 1. Introduction

This paper studies sampling efficiency in flow-based generative modeling through transport dynamics. Our premise is that few-step quality is governed by the *dynamical straightness* of the learned velocity field. The key quantity is the material derivative

$$\frac{Dv}{Dt} = \partial_t v + (v \cdot \nabla_x)v, \tag{1}$$

which measures acceleration along trajectories. When $\|Dv/Dt\|$ is large, coarse ODE integration incurs large truncation error; when it is small, trajectories approach constant-velocity transport and are easier to integrate with few function evaluations (NFEs).

[1]Independent Researcher, Mumbai, India. Correspondence to: Tauhid Khan <khantauhid313@gmail.com>.

*Presented at the 6th Muslims in Machine Learning (MusIML) Workshop at ICML 2026, Seoul, South Korea.* 

**The Eulerian–Lagrangian connection.** If the Eulerian velocity field has zero acceleration along trajectories, the induced Lagrangian flow map becomes linear: $\Phi_t(x_0) \approx x_0 + t\, v(x_0, 0)$. To achieve straight trajectories, prior works often treat this as an architectural or distillation challenge, relying on multi-stage pipelines (Liu et al., 2022; Salimans & Ho, 2022) or explicit consistency constraints (Song et al., 2023), which can introduce training instability. Instead, Iso-FM approaches flow straightening fundamentally as a continuous **regularization problem**. By directly penalizing material acceleration in the Eulerian regime, we apply a dynamic regularization that encourages straighter transport trajectories and reduced path curvature. This allows us to achieve highly competitive flow map properties while retaining the stability and simplicity of standard, single-stage FM training.

**Our contributions:**

- **Dynamical Regularization for Flow Matching.** We establish a direct theoretical link between the material derivative ($Dv/Dt$) of the learned Eulerian velocity field and the truncation error experienced during few-step ODE integration, showing that suppressing $\|Dv/Dt\|$ directly reduces Euler integration error at low NFE.

- **Jacobian-Free Optimization.** We introduce Isokinetic Flow Matching (Iso-FM), a lightweight, plug-and-play regularizer that penalizes pathwise acceleration using a self-guided finite-difference approximation, requiring only standard forward evaluations and stop-gradient targets and entirely avoiding expensive Jacobian-vector products.

- **Improved Few-Step Efficiency.** On CIFAR-10 (Krizhevsky, 2009) (DiT-S/2 (Peebles & Xie, 2023)), Iso-FM achieves a 2.9× efficiency gain at FID@2 (FID (Heusel et al., 2017): 78.82→27.13) in single-stage training, demonstrating that acceleration regularization effectively straightens generative flows without multi-stage distillation.

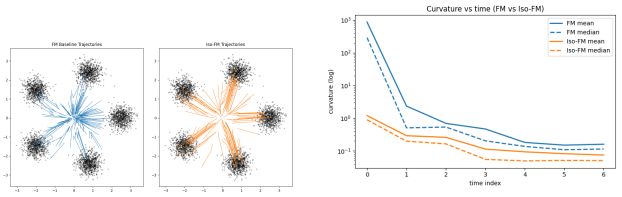


(a) Trajectory comparison (b) Curvature comparison

*Figure 1.* Low-dimensional geometric diagnostics. (a) FM trajectories (blue) exhibit stronger bending; Iso-FM trajectories (orange) are more linear. (b) Iso-FM maintains lower curvature across integration time, empirically confirming acceleration suppression.

## 2. Related Work

### 2.1. Flow Matching and Eulerian Generative Models

Generative flow models construct a time-dependent vector field $v_t(x)$ that transports a simple prior $p_0$ to a complex data distribution $p_1$ via the continuity equation:

$$\partial_t p_t + \nabla \cdot (p_t v_t) = 0. \tag{2}$$

The ideal objective matches the network $v_\theta$ to the true marginal field $u_t(x)$ generating the target path $p_t(x)$, but computing $u_t(x)$ and sampling from $p_t(x)$ directly are intractable for complex datasets.

Conditional Flow Matching (CFM) (Lipman et al., 2023) bypasses this by conditioning on individual source-target pairs $(x_0, x_1)$. Under the optimal-transport (OT) formulation, the conditional path is a linear interpolation

$$x_t = (1-t)x_0 + tx_1, \tag{3}$$

with constant conditional velocity $u_t = x_1 - x_0$. The tractable FM loss is:

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{t,x_0,x_1}[\|v_\theta(x_t,t) - (x_1 - x_0)\|^2]. \tag{4}$$

While individual conditional trajectories are perfectly straight, the network learns the *marginal* velocity at inference, which averages over all conditional fields passing through a point:

$$v(x,t) = \mathbb{E}_{x_0,x_1|x_t=x}[x_1 - x_0]. \tag{5}$$

Because independent couplings cause distinct trajectories to intersect in state space, this expectation averages conflicting velocity directions, producing strong geometric curvature in the marginal field. This curvature increases numerical integration difficulty at low NFE and acts as a major bottleneck for few-step generation that Iso-FM directly addresses.

### 2.2. Flow Straightening and Acceleration Control

*Rectified Flow* (Liu et al., 2022) iteratively refines the coupling $(x_0, x_1)$ by using samples from a pre-trained model as new data pairs, effectively straightening the flow via recursive retraining. *OT-CFM* (Tong et al., 2024) reduces curvature by matching source and target samples within a minibatch via optimal transport, creating simpler ground-truth paths. Lee et al. (Lee et al., 2023) explicitly connect solver error to path curvature and untangle intersections via $\beta$-VAE-style encoders for improved data-noise pairing. *OAT-FM* (Yue et al., 2025) minimizes $\mathbb{E}[\|\ddot{x}\|^2]$ via OT constraints as a post-training refinement step. *Shortcut Models* (Frans et al., 2025) condition on step size $d$ and enforce the self-consistency recurrence $s(x_t,t,2d) \approx \frac{1}{2}(s(x_t,t,d) + s(x_{t+d},t+d,d))$; unlike Iso-FM, they require variable-step conditioning and do not enforce local stationarity of the velocity field itself.

While these methods successfully straighten paths by modifying source-target coupling, Iso-FM achieves straightness in the *Eulerian regime* by regularizing the learned velocity field directly, without auxiliary encoders or requiring near-optimal couplings.

### 2.3. Flow Map and Consistency Methods

To bypass the computational bottleneck of sequential ODE integration, a growing body of work attempts to learn the Lagrangian solution operator $\Phi_{s\to t}$ directly. Consistency Models (Song et al., 2023) and Flow Map Matching (Boffi et al., 2025) achieve this by enforcing that distinct points sampled along the same probability flow trajectory must map to a consistent structural origin or future state. This is typically formulated via self-consistency objectives across discrete time intervals:

$$\mathcal{L}_{\text{CM}} = \mathbb{E}[\|f_\theta(x_{t+\Delta t}, t+\Delta t) - f_\theta(x_t,t)\|^2]. \tag{6}$$

While highly effective for few-step sampling, these trajectory-spanning objectives often suffer from high variance during optimization, requiring either pre-trained teacher models, specialized architectures, or careful discretization schedules.

Adapting these principles to the vector field domain, Consistency Flow Matching (Yang et al., 2024) enforces velocity consistency across time steps via Exponential Moving Average (EMA) targets. However, this relies heavily on the delicate optimization dynamics and hyperparameter tuning inherent to consistency training. Similarly targeting large-step integration, MeanFlow (Geng et al., 2025) introduces an interval-averaged velocity $U_{s\to t} = \frac{1}{t-s}\int_s^t v_\tau \, d\tau$. By enforcing the identity $v_t = U_{s\to t} + (t-s)\frac{d}{dt}U_{s\to t}$, MeanFlow mathematically bridges the instantaneous Eulerian velocity with macroscopic Lagrangian displacement to enable one-step generation.

Crucially, these explicit flow map methods require structurally modifying the learning objective to encompass multiple time steps, integral quantities, or new conditioning

parameters. In contrast, as we demonstrate mathematically in Appendix B, the discrepancy between the instantaneous velocity $v_t$ and the interval-averaged MeanFlow $U_{s\to t}$ is directly proportional to the local Eulerian acceleration $\partial_t v_t$. Thus, by suppressing material acceleration, Iso-FM reduces the leading-order discrepancy between the instantaneous velocity field and interval-averaged transport. This allows Iso-FM to promote transport dynamics favorable for large-step generation purely through local dynamic regularization, preserving the simplicity and stability of standard, single-stage Flow Matching.

**Position of Iso-FM.** We situate Iso-FM as a foundational Eulerian regularization technique. Unlike coupling methods that alter source-target pairings to straighten paths, Iso-FM directly constrains the flow dynamics by suppressing the material derivative $Dv/Dt$. Furthermore, unlike Lagrangian flow map methods that require modifying the neural architecture to condition on extra timesteps (e.g., jump intervals $\Delta t$ or target times $s$), Iso-FM encourages locally straighter transport trajectories while remaining purely Eulerian and preserving the standard $v_\theta(x,t)$ interface. Because it operates strictly via a dynamic penalty on the vector field, Iso-FM is complementary to existing coupling paradigms and can be integrated without structural overhead.

## 3. Background: Eulerian and Lagrangian Perspectives

Continuous-time generative models admit two complementary viewpoints from fluid dynamics (Boffi et al., 2025; Lipman et al., 2023). Understanding their duality is crucial for contextualizing the bottlenecks of current models and justifying the design of Iso-FM.

**Eulerian perspective (instantaneous velocity).** Flow Matching learns a time-dependent vector field $v(x,t)$ dictating the instantaneous rate of change of a sample. Generation solves the Initial Value Problem (IVP):

$$\frac{dx(t)}{dt} = v(x(t),t), \quad x(0) = x_0. \tag{7}$$

Eulerian objectives are stable and single-stage, but if the learned field contains complex geometry or crossing paths, the ODE solver requires many small steps (high NFE) to bound discretization error (Lee et al., 2023).

**Lagrangian perspective (flow map matching).** Recent frameworks such as Flow Map Matching (Boffi et al., 2025) adopt a Lagrangian perspective, directly learning a solution operator $\Phi_{s\to t}(x_s) = x_t$ that bypasses ODE integration and theoretically enables one-step sampling. However, representing the full manifold of trajectories requires conditioning on both endpoints $(s,t)$, increasing architectural complexity relative to standard FM. Furthermore, enforcing the flow map constraint $\partial_t \Phi_{s\to t} = v(\Phi_{s\to t}, t)$ requires either Jacobian-Vector Products (Boffi et al., 2025) or self-distillation with EMA teachers (Song et al., 2023), both of which introduce additional optimization complexity relative to standard FM regression.

**Bridging the gap.** The key insight driving Iso-FM: if an Eulerian velocity field has zero acceleration along its own flow lines, a particle's velocity remains constant and the corresponding Lagrangian flow maps become affine in time:

$$\Phi_{0\to t}(x_0) = x_0 + t \cdot v(x_0, 0). \tag{8}$$

This insight dictates our methodology. Rather than abandoning the stable Eulerian training paradigm to learn explicit Lagrangian flow maps, Iso-FM operates purely in the Eulerian regime. By directly suppressing pathwise acceleration, it encourages straighter Lagrangian transport trajectories while retaining single-stage FM training.

**Key observation.** If the Eulerian velocity field has *zero (or small) acceleration along trajectories*, the flow map is approximately linear:

$$\Phi_t(x_0) \approx x_0 + t\, v(x_0, 0). \tag{9}$$

Iso-FM exploits this: by straightening the Eulerian field, it implicitly induces simple Lagrangian transport without learning them explicitly.

## 4. Isokinetic Flow Matching

### 4.1. Acceleration and Sampling Error

Let $x(t)$ follow the learned ODE $\dot{x}(t) = v(x(t),t)$. The acceleration of a particle moving through the field is the *material derivative*:

$$\frac{Dv}{Dt} = \frac{\partial v}{\partial t} + (v \cdot \nabla_x) v, \tag{10}$$

decomposing into local (temporal) acceleration $\partial_t v$ and convective acceleration $(v \cdot \nabla_x)v$ due to spatial variation. A Taylor expansion of the trajectory gives:

$$x(1) = x(0) + v(x_0,0) + \tfrac{1}{2}\tfrac{Dv}{Dt}(x_0,0) + \mathcal{O}(\|\tfrac{Dv}{Dt}\|^2). \tag{11}$$

The one-step Euler error is thus governed by this acceleration: $\|x(1) - \hat{x}_1\| \approx \frac{1}{2}\|\frac{Dv}{Dt}(x_0,0)\|$ (we provide a formal bound on this local truncation error in Section C). High curvature and acceleration force numerical solvers to take many small steps.

**Conclusion:** Few-step generation requires suppressing $Dv/Dt$ along trajectories, particularly at early times $t \approx 0$ where trajectory ambiguity is greatest.

### 4.2. Self-Guided Finite-Difference Objective

Direct evaluation of $Dv/Dt$ requires expensive Jacobian-vector products. Iso-FM replaces this with a self-guided

finite-difference approximation. For a small $\varepsilon > 0$, define a lookahead step:

$$x_{t+\varepsilon} = x_t + \varepsilon\, v_\theta(x_t, t). \tag{12}$$

As we formally prove in Section C, a first-order Taylor expansion guarantees that this finite difference converges to the true material derivative:

$$\frac{v_\theta(x_{t+\varepsilon}, t+\varepsilon) - v_\theta(x_t, t)}{\varepsilon} \approx \frac{Dv_\theta}{Dt}(x_t, t). \tag{13}$$

The Iso-FM loss penalizes velocity changes along model trajectories:

$$\mathcal{L}_{\text{Iso}} = \left\| v_\theta(x_t, t) - \text{sg}[v_\theta(x_{t+\varepsilon}, t+\varepsilon)] \right\|^2, \tag{14}$$

where $\text{sg}[\cdot]$ denotes stop-gradient. This directly suppresses convective acceleration without Jacobian computations.

The total training objective seamlessly combines the standard Flow Matching regression loss with this dynamic penalty:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{FM}} + \lambda_{\text{Iso}} \mathcal{L}_{\text{Iso}}, \tag{15}$$

where $\mathcal{L}_{\text{FM}} = \|v_\theta(x_t, t) - u_t\|^2$ and $\lambda_{\text{Iso}}$ is a hyperparameter controlling the regularization strength. By optimizing $\mathcal{L}_{\text{total}}$, the network simultaneously learns the Eulerian vector field and encourages smoother and lower-curvature transport dynamics in a single stage. Furthermore, as established in Section C, Minimizing the material acceleration encourages locally straight, constant-velocity trajectories, reducing curvature in the induced transport dynamics and improving large-step numerical integration behavior.

### 4.3. Training Algorithm

Algorithm 1 summarizes the complete training step. We detail the key design choices below.

**Design choices.**

*Temporal weighting.* The factor $(1-t)^\alpha$ concentrates the isokinetic penalty at early times where the posterior over $x_1$ given $x_t$ is most diffuse, acceleration is largest, and suppression has the greatest impact on integration error (see Section D).

*Velocity normalization.* Dividing by $s = \|\text{sg}(v_{\text{curr}})\|_2 + \zeta$ (with $\zeta = 10^{-8}$ for numerical stability) normalizes the finite-difference residual by the current velocity magnitude, making $\mathcal{L}_{\text{Iso}}$ scale-invariant with respect to the overall transport speed. Without this, the regularizer would disproportionately penalize high-magnitude velocity regions.

*Log-normal $\varepsilon$ schedule.* Sampling $\varepsilon = \exp(\eta)$ with $\eta \sim \mathcal{N}(-1.5, 0.8)$ concentrates lookahead steps in the range $\varepsilon \approx 0.05$–$0.5$, matching the sub-step distances used during few-step Euler inference. The boundary constraint $\varepsilon \leq 0.999(1-t)$ ensures $t + \varepsilon \leq 1$ strictly.

**Algorithm 1** Iso-FM Training Step

**Input:** $x_1$ (data), $x_0 \sim \mathcal{N}(0, I)$, model $v_\theta$, $\lambda_{\text{FM}}$, $\lambda_{\text{Iso}}$, $\alpha$, $p_{\text{iso}}$
*// — Standard FM pass —*
$t \sim \text{LogitNormal}(0, 1)$
$x_t = (1-t)x_0 + tx_1$; $\quad u_t = x_1 - x_0$
$v_{\text{curr}} = v_\theta(x_t,\, t)$
$\mathcal{L}_{\text{FM}} = \|v_{\text{curr}} - u_t\|_2^2$
$\mathcal{L}_{\text{Iso}} \leftarrow 0$
**if** $\text{Bernoulli}(p_{\text{iso}}) = 1$ **then**
  *// — $\varepsilon$ sampling: log-normal, boundary-safe —*
  $\eta \sim \mathcal{N}(\mu_\varepsilon{=}{-1.5},\, \sigma_\varepsilon{=}0.8)$; $\quad \varepsilon \leftarrow \exp(\eta)$
  $\varepsilon \leftarrow \min\big(\varepsilon,\, 0.999(1-t)\big)$; $\quad \varepsilon \leftarrow \max(\varepsilon,\, 10^{-4})$
  *// — Self-guided lookahead (Jacobian-free) —*
  $x_{t+\varepsilon} = x_t + \varepsilon \cdot \text{sg}(v_{\text{curr}})$
  $v_{\text{next}} = v_\theta(x_{t+\varepsilon},\, t{+}\varepsilon)$ *(no grad)*
  $w = (1-t)^\alpha/\varepsilon$ *// upweights $t{\approx}0$; normalizes for step size*
  $s = \|\,\text{sg}(v_{\text{curr}})\|_2{+}10^{-6}$ *// velocity-magnitude scale; prevents shrinkage*
  $\mathcal{L}_{\text{Iso}} = \big(w \cdot \big\|\frac{v_{\text{curr}} - \text{sg}(v_{\text{next}})}{s}\big\|_1\big)_{\text{mean over batch}}$
**end if**
**return** $\mathcal{L} = \lambda_{\text{FM}}\mathcal{L}_{\text{FM}} + \lambda_{\text{Iso}}\mathcal{L}_{\text{Iso}}$

*Stochastic gate $p_{iso}$.* Applying the iso regularizer with probability $p_{\text{iso}}$ per step provides an interpolation between pure FM ($p_{\text{iso}}{=}0$) and fully regularized training ($p_{\text{iso}}{=}1$).

**Hyperparameters.** All reported experiments use $\lambda_{\text{FM}}{=}1.0$, $\lambda_{\text{Iso}}{=}4.0$, $\alpha{=}2.0$, $p_{\text{iso}}{=}1.0$. These values were set by inspection and were *not* ablated; a systematic sensitivity study is left to future work.

Iso-FM adds one extra forward pass per training step with no second-order differentiation, incurring modest wall-clock overhead ($<25\%$ in our setup).

### 4.4. OT Coupling (Optional)

For OT-enabled runs, we apply minibatch OT coupling (Tong et al., 2024) before FM interpolation. Given Gaussian samples $\{x_0^{(i)}\}_{i=1}^B$ and data $\{x_1^{(j)}\}_{j=1}^B$, we solve

$$\pi^* = \arg\min_{\pi \in S_B} \sum_{i=1}^{B} \|x_0^{(i)} - x_1^{(\pi(i))}\|_2^2 \tag{16}$$

via the Hungarian algorithm and apply FM on $x_t^{(i)} = (1-t)x_0^{(i)} + t\, x_1^{(\pi^*(i))}$. This isolates the effect of improved source-target pairing from the Iso-FM dynamical regularization term.

### 4.5. Path Geometry and Interpretation

Iso-FM does not constrain endpoints or enforce map consistency. Instead, it:

- straightens the Eulerian velocity field by suppressing $Dv/Dt$,
- reduces kinetic energy variation along trajectories,
- and encourages straighter and lower-curvature transport trajectories.

**Low-Dimensional Path Geometry Analysis.** To provide mechanistic evidence for path straightening, we include a controlled low-dimensional comparison between baseline FM and Iso-FM. Let trajectories evolve under the learned velocity field:

$$\dot{x}(t) = v_\theta(x(t), t), \qquad \ddot{x}(t) = \frac{Dv_\theta}{Dt} = \partial_t v_\theta + (v_\theta \cdot \nabla_x) v_\theta. \tag{17}$$

Iso-FM directly penalizes temporal velocity variation along model trajectories, which reduces $\|Dv_\theta/Dt\|$ and therefore suppresses path bending. For geometric quantification, we measure curvature via the proxy

$$\kappa(t) = \frac{\|\ddot{x}(t)\|}{\|\dot{x}(t)\|^2 + \varepsilon}, \tag{18}$$

where $\varepsilon > 0$ ensures numerical stability at low speed. Since $\ddot{x}(t)$ is the material acceleration, lower acceleration under Iso-FM implies lower curvature and a straighter transport map. Figure 1 confirms this: (a) FM trajectories exhibit stronger bending; (b) Iso-FM maintains consistently lower curvature across integration time, empirically supporting the theoretical objective of acceleration suppression.

Importantly, Iso-FM incurs minimal memory overhead and integrates seamlessly into standard FM training: one extra forward pass per step (<25% wall-clock overhead), no second-order differentiation, and no auxiliary encoders.

The term *isokinetic* reflects this: approximately preserved speed (kinetic energy) along trajectories; directional change is permitted, but rapid acceleration and curvature are penalized.

## 5. Experiments

All experiments use a DiT-S/2 backbone (Peebles & Xie, 2023) (∼32.5M parameters) on CIFAR-10 (Krizhevsky, 2009) (32×32), trained for 2500 epochs with AdamW (Loshchilov & Hutter, 2019) (lr=$5\times10^{-4}$), cosine schedule, EMA decay 0.9999, and bfloat16 precision. Every 250 epochs, we evaluate FID (Heusel et al., 2017) using 50,000 generated samples at NFE$\in \{1, 2, 4\}$.

*Table 1.* Best FID scores on CIFAR-10 (DiT-S/2), 2500 epochs. All pairs share architecture, training horizon, and evaluation protocol. Δ@2 is relative FID@2 improvement vs. matched FM baseline. †Cond. OT FM baseline not run due to compute budget constraints; the OT contribution is isolated and quantified in the unconditional paired rows.

| Setup | Method | FID@1↓ | FID@2↓ | FID@4↓ | Δ@2 |
|---|---|---|---|---|---|
| Cond. | FM | 245.37 | 78.82 | 27.29 | — |
| | **Iso-FM** | **83.85** | **27.13** | **15.54** | **−65.6%** |
| Cond. OT | **Iso-FM**† | **104.30** | **17.76** | **10.23** | — |
| Uncond. | FM | 327.20 | 98.29 | 47.77 | — |
| | **Iso-FM** | **143.67** | **36.92** | **25.05** | **−62.4%** |
| Uncond. OT | FM | 249.09 | 79.70 | 42.47 | — |
| | **Iso-FM** | **170.31** | **36.69** | **24.08** | **−54.0%** |

Conditional runs use CFG (Ho & Salimans, 2021) scales $\{2.5, 3.5, 4.5\}$ with label-drop probability 0.15. We report the best FID per metric across epochs. Full configuration details are in Section A.

**Scope.** We validate our theoretical findings on CIFAR-10 with a DiT-S/2 backbone. We specifically select this controlled environment because it permits the training of multiple matched baselines over 2500 epochs, ensuring that the isolated effect of Eulerian acceleration regularization is statistically clear and not confounded by massive compute scale. Extending these principles to larger-scale datasets and models is a natural next step for future work.

### 5.1. Main Results

Table 1 reports best-observed FID scores across all seven runs. Iso-FM consistently improves low-NFE FID versus matched FM baselines in every setting.

The headline result is a 2.9× efficiency gain at conditional FID@2 (78.82→27.13, −65.6%). Across all three paired settings the FID@2 reduction ranges from 54% to 66%, confirming the consistency of the effect. Adding OT coupling to Iso-FM in the conditional setting yields the best single scores: **FID@2=17.76** and **FID@4=10.23**; since the matched Cond. OT FM baseline was not run, the OT contribution is quantified cleanly in the unconditional paired rows above.

### 5.2. Training Dynamics

Figures 3 and 4 show FID trajectories over training. Iso-FM converges to lower FID@2 and FID@4 earlier in training compared to FM baselines, suggesting that acceleration regularization also acts as an inductive bias favoring straighter flows from early epochs.

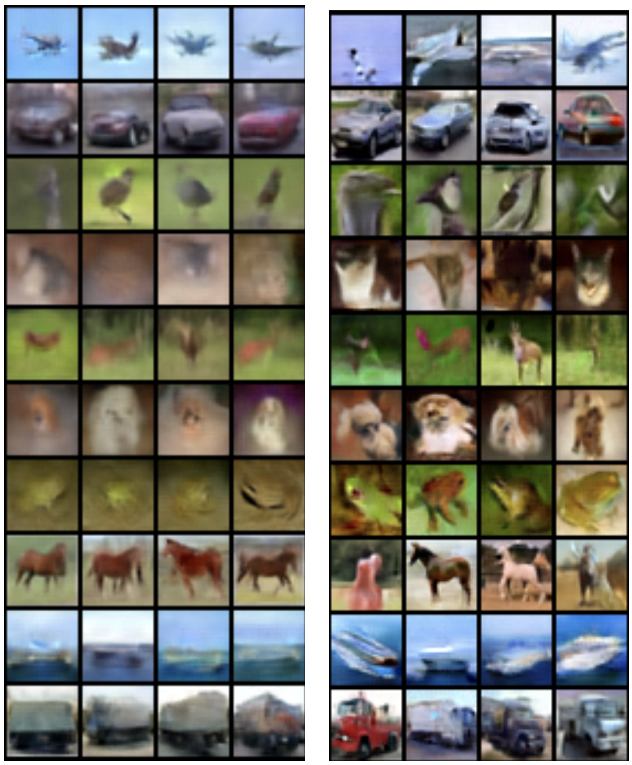

FM (FID@2 = 78.82) Iso-FM (FID@2 = 27.13)

*Figure 2.* Conditional CIFAR-10 samples generated with **only 2 NFE**. Left: standard FM (FID = 78.82). Right: Iso-FM (FID = 27.13, −65.6%). Iso-FM produces markedly sharper class structure and finer detail at the same compute budget, illustrating the qualitative impact of acceleration regularization.

### 5.3. Analysis of Results

Iso-FM consistently improves few-step FID across all settings. The headline conditional non-OT result at FID@2 is 78.82 → 27.13, a **65.6% FID reduction** (2.9× efficiency gain). In unconditional non-OT runs, FID@2 improves by 62.4% (98.29 → 36.92). In OT runs, FID@2 improves 54.0% (79.70 → 36.69) and FID@4 by 43.3% (42.47 → 24.08) vs. the OT-FM baseline. The best configuration (Cond. OT+Iso-FM) reaches FID@4=10.23.

These gains are achieved with <25% wall-clock overhead (one extra forward pass per step) and no second-order differentiation, making Iso-FM a practical add-on to any single-stage FM pipeline.

## 6. Conclusion

Iso-FM regularizes the material derivative of the learned velocity field in a Jacobian-free manner, integrated as a plug-and-play term in standard single-stage FM training. Under matched DiT-S/2 CIFAR-10 settings, it consistently improves low-NFE performance, including a 2.9× efficiency gain at conditional FID@2. These results indicate that path-

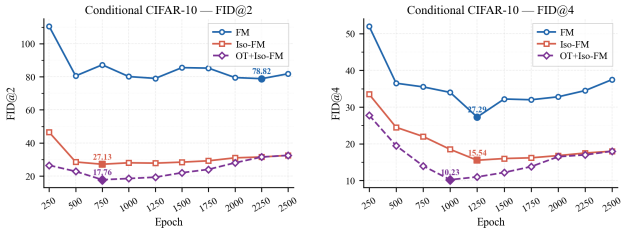


(a) Conditional FID@2 (b) Conditional FID@4

*Figure 3.* Conditional CIFAR-10 training dynamics: FID@2 and FID@4 vs. epoch for FM, Iso-FM, and OT+Iso-FM.

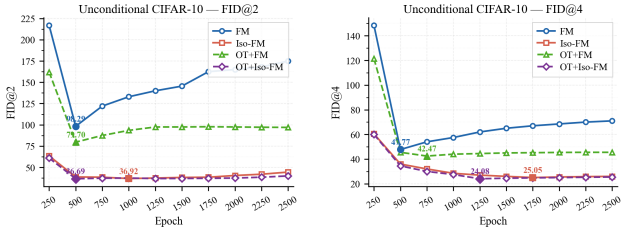


(a) Unconditional FID@2 (b) Unconditional FID@4

*Figure 4.* Unconditional CIFAR-10 training dynamics: FID@2 and FID@4 vs. epoch for FM, Iso-FM, OT+FM, and OT+Iso-FM.

wise acceleration regularization shifts the quality–compute trade-off favorably without requiring multi-stage distillation or architectural changes.

**Claim scope.** Iso-FM is a methodological contribution focused on dynamics-aware efficiency, not an absolute state-of-the-art benchmark claim. While we observe strong relative gains at low NFE under controlled single-stage FM training, we do not claim superiority over heavily scaled or multi-stage distilled pipelines. The strength of the result lies in isolating the effect of Eulerian acceleration regularization under matched conditions across seven experimental settings.

**Straightness vs. generative variability.** A natural concern: does penalizing acceleration restrict sample diversity? As shown in Section D, the marginal acceleration $Dv/Dt$ is inextricably tied to the conditional velocity variance $\Sigma(x,t) = \mathbb{E}_{z|x_t}[(u-v)(u-v)^T]$, which encodes multimodality. A hard constraint

$$Dv/Dt \equiv 0$$

would require the covariance-induced forcing term

$$\nabla_x \cdot (p\Sigma)$$

to vanish, potentially restricting the multimodal transport structure needed to represent complex data distributions. Because Iso-FM applies a *soft* regularizer ($\lambda_{\text{Iso}}\mathcal{L}_{\text{Iso}}$), the primary FM objective preserves $\Sigma(x,t)$ and sample diversity. The regularizer suppresses unnecessary convective accelerations (twisted or crossing trajectories) that inflate truncation

errors, while allowing the macroscopic acceleration needed for samples to commit to specific modes. Our FID scores confirm that Iso-FM straightens transport geometry without sacrificing generative diversity.

**Limitations and future work.** Several directions remain open.

*Scale.* All experiments use CIFAR-10 at $32\times32$ with a 32M-parameter backbone. Whether the gains in low-NFE quality transfer to larger datasets (e.g., ImageNet $256\times256$) and larger DiT variants is an important open question. The compute cost of the extra forward pass grows linearly with model size, so efficiency in the overhead itself may require future engineering.

*Hyperparameter sensitivity.* The isokinetic weight $\lambda_{\text{Iso}}$, the temporal exponent $\alpha$, and the log-normal $\varepsilon$ schedule were set by inspection and not ablated. A principled sensitivity study, or an adaptive scheme based on local velocity variance, is needed before claiming robustness.

*Hard acceleration lower bound.* As proven in Section D, the canonical marginal Flow Matching velocity field generally exhibits nonzero acceleration under multimodal conditional uncertainty; Iso-FM suppresses *unnecessary* curvature but cannot eliminate the residual acceleration associated with multimodal transport structure for the model to commit to distinct data modes.

*Hybrid Eulerian-Lagrangian training.* A natural next step is combining Iso-FM with an explicit flow-map head: Iso-FM would reduce stiffness of the Eulerian field while the map head could amortize any remaining integration error, potentially enabling true one-step generation without multi-stage distillation. Preliminary low-dimensional experiments suggest this combination is stable, but a systematic study is left for future work.

## Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning, specifically by improving the sampling efficiency of continuous-time generative models. On the positive side, this improved efficiency can lower the energy consumption and carbon footprint of model deployment, while democratizing access to generative tools for researchers with limited compute. However, we acknowledge that accelerating generative models simultaneously lowers the barrier for generating synthetic media at scale, which can be misused for disinformation. We encourage future deployments of highly efficient generative models to be paired with robust provenance tooling and standard safety filters.

## References


Boffi, N. M., Albergo, M. S., and Vanden-Eijnden, E. Flow map matching with stochastic interpolants: A mathematical framework for consistency models, 2025. URL https://arxiv.org/abs/2406.07507. arXiv:2406.07507.

Frans, K., Hafner, D., Levine, S., and Abbeel, P. One step diffusion via shortcut models, 2025. URL https://arxiv.org/abs/2410.12557. arXiv:2410.12557.

Geng, Z., Deng, M., Bai, X., Kolter, J. Z., and He, K. Mean flows for one-step generative modeling, 2025. URL https://arxiv.org/abs/2505.13447. arXiv:2505.13447.

Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, 2017. URL https://arxiv.org/abs/1706.08500.

Ho, J. and Salimans, T. Classifier-free diffusion guidance, 2021. URL https://arxiv.org/abs/2207.12598. arXiv:2207.12598.

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

Lee, S., Kim, B., and Ye, J. C. Minimizing trajectory curvature of ODE-based generative models, 2023. URL https://arxiv.org/abs/2301.12003. arXiv:2301.12003.

Lipman, Y., Chen, R. T. Q., Ben-Hamu, H., Nickel, M., and Le, M. Flow matching for generative modeling, 2023. URL https://arxiv.org/abs/2210.02747. arXiv:2210.02747.

Liu, X., Gong, C., and Liu, Q. Flow straight and fast: Learning to generate and transfer data with rectified flow, 2022. URL https://arxiv.org/abs/2209.03003. arXiv:2209.03003.

Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2019. URL https://arxiv.org/abs/1711.05101.

Peebles, W. and Xie, S. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4195–4205, 2023. URL https://arxiv.org/abs/2212.09748.

Salimans, T. and Ho, J. Progressive distillation for fast sampling of diffusion models, 2022. URL https://arxiv.org/abs/2202.00512.

Song, Y., Dhariwal, P., Chen, M., and Sutskever, I. Consistency models, 2023. URL https://arxiv.org/abs/2303.01469. arXiv:2303.01469.

Tong, A., Fatras, K., Malkin, N., Huguet, G., Zhang, Y., Rector-Brooks, J., Wolf, G., and Bengio, Y. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*, 2024.

Yang, L., Zhang, Z., Zhang, Z., Liu, X., Xu, M., Zhang, W., Meng, C., Ermon, S., and Cui, B. Consistency flow matching: Defining straight flows with velocity consistency, 2024. URL https://arxiv.org/abs/2407.02398. arXiv:2407.02398.

Yue, A., Dong, A., and Xu, H. OAT-FM: Optimal acceleration transport for improved flow matching, 2025. URL https://arxiv.org/abs/2509.24936. arXiv:2509.24936.

# A. Training Configuration

*Table 2.* Complete training configuration for CIFAR-10 experiments.

| Item | Value |
|---|---|
| Backbone | DiT-S/2 (Peebles & Xie, 2023) Transformer (depth=12, hidden dim=384, heads=6, patch size=2, MLP ratio=4.0). |
| Trainable parameters | 32,485,260 (conditional, num_classes=10); 32,481,804 (unconditional, num_classes=1). |
| Data / resolution | CIFAR-10 (Krizhevsky, 2009) train split, RGB 32×32, random horizontal flip augmentation. |
| Batch size / epochs | 256 / 2500. |
| Optimizer | AdamW (Loshchilov & Hutter, 2019) (lr=$5\times10^{-4}$, weight decay= $10^{-4}$). |
| Learning-rate schedule | Cosine annealing with $T_{\max}$=2500, $\eta_{\min}$=0.1×lr. |
| EMA / precision / clipping | EMA decay=0.9999, bfloat16 mixed precision, gradient clipping at global norm 1.0. |
| Sampling of $t$ | Logit-normal sampling ($\mu$=0.0, $\sigma$=1.0). |
| Loss weights | $\lambda_{\text{FM}}$=1.0 for all runs; FM baseline uses $\lambda_{\text{Iso}}$=0.0, $p_{\text{iso}}$=0.0; Iso-FM uses $\lambda_{\text{Iso}}$=4.0, $p_{\text{iso}}$=1.0, $\alpha$=2.0. |
| CFG settings (Ho & Salimans, 2021) | Conditional: label-drop probability 0.15 and eval CFG scales $\{2.5, 3.5, 4.5\}$; Unconditional: CFG scale=1.0. |
| OT setting | When enabled, minibatch coupling uses Hungarian assignment on squared Euclidean cost before interpolation. |
| Evaluation protocol | Every 250 epochs, 50,000 generated samples, FID (Heusel et al., 2017) at NFE$\in \{1, 2, 4\}$. |

# B. Connection to Mean Velocity Formulations

MeanFlow-style formulations (Geng et al., 2025) introduce an interval-averaged velocity to bypass sequential integration. In this section, we establish a rigorous mathematical connection between this averaged displacement and the material acceleration penalized by Iso-FM.

**Definition B.1** (MeanFlow Interval-Averaged Velocity)**.** Let $v_t(x)$ denote the learned velocity field, and let $\phi_{s\to t}(x)$ denote the flow map that satisfies the following:

$$\frac{d}{dt}\phi_{s\to t}(x) = v_t(\phi_{s\to t}(x)), \qquad \phi_{s\to s}(x) = x. \tag{19}$$

The MeanFlow interval-averaged velocity is defined as the macroscopic displacement divided by the time interval:

$$U_{s\to t}(x_s) = \frac{\phi_{s\to t}(x_s) - x_s}{t-s} = \frac{1}{t-s}\int_s^t v_\tau(\phi_{s\to\tau}(x_s))\, d\tau. \tag{20}$$

We emphasize that $U_{s\to t}(x_s)$ is evaluated at the *fixed* base point $x_s$, while the integrand $v_\tau(\phi_{s\to\tau}(x_s))$ tracks the *trajectory* emanating from $x_s$.

Using Theorem B.1, we can derive the exact discrepancy between the instantaneous velocity and the required mean displacement.

**Proposition B.2** (Fundamental Discrepancy Identity)**.** *The discrepancy between the instantaneous velocity $v_t$ evaluated at the target state $x_t = \phi_{s\to t}(x_s)$ and the interval-averaged velocity $U_{s\to t}$ evaluated at the source state $x_s$ is proportional to the material acceleration to leading order:*

$$v_t(x_t) - U_{s\to t}(x_s) = \frac{t-s}{2}\frac{Dv}{Dt}(x_t) + \mathcal{O}\big((t-s)^2\big)\,. \tag{21}$$

*Proof.* To connect $U_{s\to t}$ to quantities at time $t$, we perform a Taylor expansion of the integrand $v_\tau(\phi_{s\to\tau}(x_s))$ *along the trajectory* (i.e., in the Lagrangian frame). Define the trajectory-evaluated velocity:

$$\tilde{v}(\tau) := v_\tau(\phi_{s\to\tau}(x_s))\,. \tag{22}$$

Its total time derivative is exactly the *material derivative* of $v$:

$$\frac{d\tilde{v}}{d\tau} = \partial_\tau v_\tau(\phi_{s\to\tau}(x_s)) + (v_\tau \cdot \nabla_x v_\tau)(\phi_{s\to\tau}(x_s)) = \frac{Dv}{D\tau}(\phi_{s\to\tau}(x_s))\,. \tag{23}$$

Applying a first-order Taylor expansion of $\tilde{v}(\tau)$ around $\tau = t$ gives:

$$\tilde{v}(\tau) = \tilde{v}(t) + (\tau - t)\frac{d\tilde{v}}{d\tau}\Big|_{\tau=t} + R_2(\tau, t), \tag{24}$$

where the remainder $R_2(\tau,t) = \int_t^\tau (\tau-\sigma)\frac{d^2\tilde{v}}{d\sigma^2}\big|_\sigma\, d\sigma$ satisfies $|R_2(\tau,t)| = \mathcal{O}\big((\tau-t)^2\big)$.

Substituting Equation (24) into the integral definition of Equation (20) and noting that $\tilde{v}(t) = v_t(x_t)$:

$$\begin{aligned} U_{s\to t}(x_s) &= \frac{1}{t-s}\int_s^t \left[v_t(x_t) + (\tau - t)\frac{Dv}{Dt}(x_t) + R_2(\tau,t)\right] d\tau \\ &= v_t(x_t) + \frac{1}{t-s}\cdot\frac{(\tau-t)^2}{2}\frac{Dv}{Dt}(x_t)\Big|_{\tau=s}^{\tau=t} + \frac{1}{t-s}\int_s^t R_2(\tau,t)\,d\tau \\ &= v_t(x_t) - \frac{(t-s)}{2}\frac{Dv}{Dt}(x_t) + \mathcal{O}\big((t-s)^2\big)\,, \end{aligned} \tag{25}$$

where we used the fact that $\frac{1}{t-s}\int_s^t(\tau - t)\,d\tau = -\frac{(t-s)}{2}$ and the remainder satisfies $\frac{1}{t-s}\int_s^t |R_2(\tau,t)|\,d\tau = \mathcal{O}((t-s)^2)$. Rearranging Equation (25) yields the exact (to first order) discrepancy identity. □

*Remark* B.3 (Eulerian vs. Lagrangian Evaluation Points). Theorem B.2 compares $v_t$ evaluated at $x_t = \phi_{s\to t}(x_s)$ with $U_{s\to t}$ evaluated at the base point $x_s$. For large step sizes $(t-s)$, the displacement $x_t - x_s \approx (t-s)v_t$ is non-negligible, meaning the two sides are not directly comparable at the same spatial location. The identity is therefore most naturally interpreted in the Lagrangian frame along the trajectory.

*Remark* B.4 (Incompleteness of Eulerian-Only Proofs). A naive Taylor expansion of $v_\tau(x)$ at a *fixed* spatial point $x$ uses only the partial derivative $\partial_t v_t$ and misses the convective term $(v\cdot\nabla_x)v$. This omits a term of the same order as the Eulerian acceleration in general flow fields and conflates the Eulerian and Lagrangian frames. Equation (23) confirms that the correct derivative to use is the full material derivative $Dv/Dt$.

*Remark* B.5 (Implications for Iso-FM). Theorem B.2 provides a precise theoretical justification for Iso-FM's regularization strategy. If the material acceleration is globally suppressed ($\|Dv/Dt\| \to 0$), then:

$$v_t(x_t) \approx U_{s\to t}(x_s) \qquad \forall\, 0 \le s < t \le 1. \tag{26}$$

This demonstrates that the instantaneous velocity field becomes a consistent approximation to the mean displacement at *any* scale $(t-s)$, which is the exact condition required for accurate large-step ODE integration. Crucially, while MeanFlow achieves this by explicitly modifying the regression target to span multiple timesteps, Iso-FM enforces the same necessary condition *locally* through material acceleration regularization, thus preserving the original single-stage Flow Matching training objective.

## C. Theoretical Analysis of Isokinetic Flow Matching

In this section, we formalize the relationship between Eulerian material acceleration, numerical integration error, and the convergence of our finite-difference surrogate.

### C.1. Formal Bounds on Local Truncation Error

**Definition C.1.** Let $v(x,t) : \mathbb{R}^d \times [0,1] \to \mathbb{R}^d$ be a sufficiently smooth time-dependent vector field. For a particle trajectory $x(t)$ governed by $dx(t)/dt = v(x(t),t)$, the material (Lagrangian) acceleration is:

$$\frac{Dv}{Dt}(x,t) = \frac{\partial v}{\partial t}(x,t) + \big(v(x,t)\cdot\nabla_x\big)\,v(x,t). \tag{27}$$

**Proposition C.2** (Local Truncation Error for Forward Euler)**.** *Let $v \in C^2(\mathbb{R}^d \times [0,1])$ and suppose the material acceleration is bounded on a compact set $K \subset \mathbb{R}^d \times [0,1]$ by*

$$M_K \;=\; \sup_{(x,t)\in K} \left\| \frac{Dv}{Dt}(x,t) \right\| \;<\; \infty.$$

*Let $x(t)$ be the exact trajectory with $x(t_0) = x_0$, remaining in $K$ on $[t_0,\, t_0 + h]$. The local truncation error of a single Forward Euler step is:*

$$\big\| x(t_0 + h) - (x_0 + h\, v(x_0, t_0)) \big\| \;\le\; \frac{h^2}{2}\, M_K. \tag{28}$$

*Proof.* By Taylor's theorem applied to $x(t)$ around $t_0$:

$$x(t_0 + h) = x(t_0) + h\, \dot{x}(t_0) + \frac{h^2}{2}\, \ddot{x}(\xi) \quad \text{for some } \xi \in (t_0,\, t_0 + h).$$

Since $\dot{x}(t) = v(x(t), t)$, differentiating again gives $\ddot{x}(t) = Dv/Dt\,(x(t), t)$. Substituting the Euler approximation $\hat{x}_1 = x_0 + h\, v(x_0, t_0)$:

$$x(t_0 + h) - \hat{x}_1 = \frac{h^2}{2}\, \frac{Dv}{Dt}(x(\xi), \xi).$$

Taking norms and using the bound $M_K$ (valid since $(x(\xi), \xi) \in K$) completes the proof. □

*Remark* C.3. For one-step generation ($h = 1$), the bound $M_K/2$ is only informative when $M_K$ is small. While driving $\|Dv/Dt\| \to 0$ via regularization reduces this bound, two caveats apply. First, $M_K$ must be controlled on the *actual support* of the data, not globally; the bound is vacuous if $M_K$ is large. Second, this error is separate from the approximation error of $v_\theta$ to the true probability-flow velocity; both sources of error contribute to the final generation quality.

### C.2. Approximation Error of the Finite-Difference Surrogate

**Definition C.4.** Define the exact acceleration penalty as $\mathcal{L}^*(x,t) = \|\frac{Dv_\theta}{Dt}(x,t)\|^2$. For a lookahead step $\varepsilon > 0$, the finite-difference surrogate is:

$$\tilde{\mathcal{L}}_\varepsilon(x,t) = \frac{1}{\varepsilon^2} \left\| v_\theta\big(x + \varepsilon\, v_\theta(x,t),\; t + \varepsilon\big) - v_\theta(x,t) \right\|^2. \tag{29}$$

*Remark* C.5 (Degeneracy Warning). $\tilde{\mathcal{L}}_\varepsilon \equiv 0$ is achieved by any constant velocity field $v_\theta \equiv c$, including the trivial $c = 0$. In practice, this surrogate *must* be combined with a flow-matching loss that enforces the boundary conditions $v_\theta(x, 0) \approx v_{\text{target}}$; without it, minimizing $\tilde{\mathcal{L}}_\varepsilon$ alone yields degenerate solutions.

**Proposition C.6** (Pointwise Surrogate Convergence)**.** *For a fixed $v_\theta \in C^2(\mathbb{R}^d \times [0, 1+\delta])$, where $\delta > 0$ smoothly accommodates the maximum lookahead step $\varepsilon \le \delta$, and any fixed $(x,t) \in \mathbb{R}^d \times [0,1]$,*

$$\tilde{\mathcal{L}}_\varepsilon(x,t) = \mathcal{L}^*(x,t) + \mathcal{O}(\varepsilon) \quad \textit{as } \varepsilon \to 0, \tag{30}$$

*where the implicit constant depends on the $C^2$ norm of $v_\theta$ on $\mathbb{R}^d \times [0, 1+\delta]$. The base point $(x,t)$ is restricted to $\mathbb{R}^d \times [0,1]$, while the evaluation point $(x + \varepsilon v_\theta(x,t),\, t + \varepsilon)$ lies in $\mathbb{R}^d \times [0, 1+\delta]$, which is within the extended smooth domain by construction.*

*Proof.* Expanding $v_\theta(x + \varepsilon v_\theta(x,t),\, t + \varepsilon)$ by a first-order Taylor expansion around $(x,t)$:

$$\begin{aligned} v_\theta(x + \varepsilon v_\theta,\, t + \varepsilon) &= v_\theta(x,t) + \varepsilon\, \partial_t v_\theta(x,t) + \varepsilon \big(\nabla_x v_\theta(x,t)\big) v_\theta(x,t) + \mathcal{O}(\varepsilon^2) \\ &= v_\theta(x,t) + \varepsilon\, \frac{Dv_\theta}{Dt}(x,t) + \mathcal{O}(\varepsilon^2). \end{aligned} \tag{31}$$

Subtracting $v_\theta(x,t)$, dividing by $\varepsilon$, and squaring:

$$\frac{1}{\varepsilon^2} \left\| \varepsilon\, \frac{Dv_\theta}{Dt} + \mathcal{O}(\varepsilon^2) \right\|^2 = \left\| \frac{Dv_\theta}{Dt} \right\|^2 + \mathcal{O}(\varepsilon) = \mathcal{L}^*(x,t) + \mathcal{O}(\varepsilon). \tag{32}$$

□

### C.3. Geometry of Isokinetic Flows

**Lemma C.7** (Straight-Line Trajectories)**.** *If $v \in C^1(\mathbb{R}^d \times [0,1])$ satisfies $Dv/Dt \equiv 0$, then every Lagrangian particle trajectory is a straight line traversed at constant speed:*

$$x(t) = x_0 + t\, v(x_0, 0), \qquad t \in [0,1]. \tag{33}$$

*Consequently, the flow map $\Phi_t(x_0) = x_0 + t\, v(x_0, 0)$ is **affine in** $t$ but is generally **nonlinear in** $x_0$.*

*Proof.* $Dv/Dt \equiv 0$ means $\frac{d}{dt} v(x(t), t) = 0$ along any trajectory, so $v(x(t), t) = v(x_0, 0) =: c$ is constant. Integrating $\dot{x} = c$ with $x(0) = x_0$ gives $x(t) = x_0 + tc$. Since $v(x_0, 0)$ is generally a nonlinear function of $x_0$, the map $x_0 \mapsto \Phi_t(x_0)$ inherits this nonlinearity. □

*Remark* C.8. Straight-line particle paths are a *necessary* condition for quadratic-cost Optimal Transport (OT) geodesics, but they are not *sufficient*. OT additionally requires the transport map $x_0 \mapsto \Phi_1(x_0)$ to be the gradient of a convex function (Brenier's theorem), which is an independent, global condition on the coupling between source and target. Isokinetic regularization encourages straight-line paths and thereby removes one obstruction to OT-optimality, but a complete OT characterization would require further assumptions on the flow matching objective and the resulting coupling.

## D. The Fundamental Limit of Flow Straightening

While Iso-FM successfully reduces the material acceleration $\frac{Dv}{Dt}$ via mini-batch curvature regularization, we demonstrate that the canonical target of the flow matching objective possesses an intrinsic acceleration. This curvature is not a failure of the regularization scheme, but a fundamental consequence of the source-target coupling. We establish this by deriving an exact identity connecting the marginal acceleration to the conditional velocity variance, framing the effect of conditional uncertainty as a Reynolds-stress forcing term in the marginal dynamics.

**Definition D.1.** Let $z$ be a generalized conditioning variable (e.g., a specific target $z = x_1$, or a coupled source-target pair $z = (x_0, x_1)$). Let $p(x, t|z)$ and $u(x, t|z)$ be the conditional probability density and conditional velocity field respectively. The marginal density $p(x, t)$ and marginal velocity $v(x, t)$ are defined by marginalization over the coupling distribution $q(z)$:

$$p(x,t) = \int p(x,t|z)\, q(z)\, dz, \tag{34}$$

$$v(x,t) = \frac{1}{p(x,t)} \int u(x,t|z)\, p(x,t|z)\, q(z)\, dz = \mathbb{E}_{z|x_t=x}[u]. \tag{35}$$

**Assumption D.2.** We assume the conditional velocity field $u(\cdot|z)$ is at least $C^1$ in space and time, and that the coupling $q(z)$ and conditional density $p(\cdot|z)$ possess sufficient regularity such that the marginal velocity field $v(x, t)$ is also $C^1$. Furthermore, we assume finite second moments, $\mathbb{E}_{z|x_t=x}[\|u\|^2] < \infty$, ensuring the conditional covariance is well-defined. By construction in Optimal Transport Flow Matching and Rectified Flow, the conditional paths are straight lines, so their material acceleration vanishes:

$$\frac{Du}{Dt} = \partial_t u + (u \cdot \nabla_x)\, u = 0. \tag{36}$$

Both the conditional and marginal fields satisfy their respective continuity equations: $\partial_t p(\cdot|z) + \nabla_x \cdot (p(\cdot|z)\, u) = 0$ and $\partial_t p + \nabla_x \cdot (p\, v) = 0$. We additionally assume that $\partial_t$ and $\nabla_x \cdot$ may be exchanged with the $z$-integral wherever required below; a sufficient condition is that $p(\cdot|z)\, u$ and $p(\cdot|z)\, uu^T$ are dominated by a $q$-integrable function, uniformly in $(x, t)$.

**Theorem D.3** (Reynolds Decomposition of Marginal Acceleration)**.** *Under Theorems D.1 and D.2, the material derivative of the marginal velocity field satisfies, $p$-almost everywhere:*

$$\frac{Dv}{Dt}(x,t) = -\frac{1}{p(x,t)}\, \nabla_x \cdot \big(p(x,t)\, \Sigma(x,t)\big), \tag{37}$$

*where $\Sigma(x,t) = \mathbb{E}_{z|x_t=x}\big[(u-v)(u-v)^T\big]$ is the conditional velocity covariance tensor. Consequently, $Dv/Dt \equiv 0$ holds $p$-a.e. if and only if $\nabla_x \cdot (p\, \Sigma) \equiv 0$ $p$-a.e.*

*Proof.* We evaluate $\partial_t(pv)$ and then isolate $Dv/Dt$ via the product rule.

**Step 1: Time derivative of marginal momentum.** Exchanging $\partial_t$ with the $z$-integral (justified by the domination condition on $p(\cdot|z)\,u$) and applying the product rule:

$$\partial_t(pv) = \int \big(u\,\partial_t p(\cdot|z) + p(\cdot|z)\,\partial_t u\big)\,q(z)\,dz. \tag{38}$$

Substituting the conditional continuity equation $\partial_t p(\cdot|z) = -\nabla_x \cdot (p(\cdot|z)\,u)$:

$$\partial_t(pv) = \int \Big(-u\,\nabla_x \cdot (p(\cdot|z)\,u) + p(\cdot|z)\,\partial_t u\Big)\,q(z)\,dz. \tag{39}$$

**Step 2: Vector-calculus identity.** The pointwise identity (verified componentwise) $\nabla_x \cdot (p\,uu^T) = u\,\nabla_x \cdot (p\,u) + p\,(u \cdot \nabla_x)u$ allows us to write $u\,\nabla_x \cdot (p(\cdot|z)\,u) = \nabla_x \cdot (p(\cdot|z)\,uu^T) - p(\cdot|z)(u \cdot \nabla_x)u$, giving:

$$\partial_t(pv) = \int \Big(-\nabla_x \cdot (p(\cdot|z)\,uu^T) + p(\cdot|z)\,\underbrace{\big[\partial_t u + (u \cdot \nabla_x)u\big]}_{Du/Dt\,=\,0}\Big)\,q(z)\,dz. \tag{40}$$

Exchanging $\nabla_x \cdot$ with the $z$-integral (justified by the domination condition on $p(\cdot|z)\,uu^T$):

$$\partial_t(pv) = -\nabla_x \cdot \big(p(x,t)\,\mathbb{E}_{z|x_t=x}[uu^T]\big). \tag{41}$$

**Step 3: Reynolds decomposition.** The identity $\mathbb{E}[uu^T] = vv^T + \Sigma$ gives:

$$\partial_t(pv) = -\nabla_x \cdot (p\,vv^T) - \nabla_x \cdot (p\,\Sigma). \tag{42}$$

**Step 4: Isolating the material derivative.** Expanding via the product rule and grouping:

$$v\,\underbrace{\big[\partial_t p + \nabla_x \cdot (pv)\big]}_{=\,0} + p\,\underbrace{\big[\partial_t v + (v \cdot \nabla_x)v\big]}_{=\,Dv/Dt} = -\nabla_x \cdot (p\,\Sigma). \tag{43}$$

The first bracket vanishes by the marginal continuity equation. Dividing by $p(x,t)$ for all $(x,t)$ such that $p(x,t) > 0$ yields the stated identity. □

**Proposition D.4** (Projection Property of Flow Matching)**.** *The standard flow matching objective $\mathcal{L}_{\mathrm{FM}}(v_\theta) = \mathbb{E}_{t,x,z}\big[\|v_\theta(x,t) - u(x,t|z)\|^2\big]$ is uniquely minimized (up to sets of measure zero) by the conditional expectation. By conditioning on $(x,t)$ and applying the orthogonality principle of conditional expectation, the loss decomposes for any parameterized vector field $v_\theta$ as:*

$$\mathcal{L}_{\mathrm{FM}}(v_\theta) = \mathcal{L}_{\mathrm{FM}}(v^*) + \mathbb{E}_{t,x}\big[\|v_\theta(x,t) - v^*(x,t)\|^2\big], \tag{44}$$

*where $v^*(x,t) = v(x,t) = \mathbb{E}_{z|x_t=x}[u(x,t|z)]$ is the canonical marginal velocity field. Consequently, any deviation $v_\theta \neq v^*$ on a set of non-zero measure strictly increases the flow matching loss.*

**Corollary D.5** (Intrinsic Acceleration of the Canonical Target)**.** *Let $v^*(x,t)$ denote the canonical population regression target of the flow matching objective. By Theorem D.3, its intrinsic material acceleration is exactly:*

$$\frac{Dv^*}{Dt}(x,t) = -\frac{1}{p(x,t)}\,\nabla_x \cdot \big(p(x,t)\,\Sigma(x,t)\big). \tag{45}$$

*Consequently, the expected squared acceleration of this canonical target field is:*

$$\mathbb{E}_{x,t}\left[\left\|\frac{Dv^*}{Dt}\right\|^2\right] = \mathbb{E}_{x,t}\left[\left\|\frac{1}{p}\,\nabla_x \cdot (p\,\Sigma)\right\|^2\right]. \tag{46}$$

*For any learned field $v_\theta$ trained under the combined Iso-FM objective, any velocity field achieving lower acceleration than the canonical target must differ from $v^*$ on a set of non-zero measure, and therefore incurs a strictly higher population flow matching loss by Theorem D.4.*

*Remark* D.6 (What Iso-FM Regularization Achieves). Theorem D.5 precisely characterizes the role of Iso-FM's mini-batch curvature regularization. The learned field $v_\theta$ in a standard flow matching model can exhibit two distinct sources of curvature:

1. **Intrinsic canonical curvature**, given by Equation (45), which is governed by the source-target coupling and cannot be eliminated without deviating from the canonical FM regression target.
2. **Excess network-induced curvature**, arising from the network parameterization and finite training dynamics, which is not required by the flow matching objective and can in principle be suppressed.

Iso-FM's regularization term directly targets and suppresses the second source. As training progresses and $v_\theta \to v^*$, the residual acceleration converges to the intrinsic level defined in Equation (46). The conditional uncertainty intrinsic to stochastic or mini-batch OT couplings induces a Reynolds-stress forcing term ($\nabla_x \cdot (p\,\Sigma)$). Therefore, rather than achieving perfectly straight marginal paths, the theoretical optimum of the flow matching objective exhibits an intrinsic curvature governed by the spatial divergence of the conditional variance.

## E. Limitations and Open Problems

While Isokinetic Flow Matching (Iso-FM) provides a principled and lightweight mechanism for reducing trajectory curvature and enabling efficient few-step sampling, several limitations and open questions remain.

### E.1. Limits of One-Step Generation

Iso-FM substantially reduces material acceleration of the learned velocity field, which in turn lowers discretization error for coarse numerical solvers. However, this alone does not guarantee high-fidelity *one-step* generation in complex, multimodal data distributions.

In particular, when the conditional posterior $p(x_1 \mid x_0)$ is strongly multimodal at early times, the optimal marginal velocity remains an average over competing directions. While Iso-FM suppresses rapid changes in this average along trajectories, it does not resolve the fundamental ambiguity at $t = 0$. Empirically, this manifests as a performance gap between 1-step and 2-step sampling that persists even under strong regularization.

This limitation is intrinsic to instantaneous velocity regression and suggests that true one-step generation may require either: (i) explicit flow-map learning, (ii) additional symmetry-breaking mechanisms, or (iii) auxiliary conditioning information.

### E.2. Hyperparameter Sensitivity

The effectiveness of Iso-FM depends on several hyperparameters, including the isokinetic weight $\lambda$, the time-weighting exponent $\alpha$, and the distribution used to sample the lookahead step $\varepsilon$.

While we find that log-normal or Beta-distributed $\varepsilon$ schedules are robust across tasks, a principled method for selecting these parameters remains an open problem. Adaptive schemes based on local curvature, posterior entropy, or velocity variance may offer a more automatic alternative.

### E.3. Theoretical Generalization Bounds

While we provide local Taylor-based arguments linking Iso-FM to reduced Euler integration error, global generalization guarantees for long-time transport remain open.

In particular, bounding the accumulated error of few-step solvers as a function of the expected material acceleration along learned trajectories is an important theoretical question that may connect Iso-FM to recent advances in neural ODE stability analysis.

## F. Additional Generated Samples

*Figure 5.* Conditional samples at epoch: FM baseline (left) vs. Iso-FM (right).

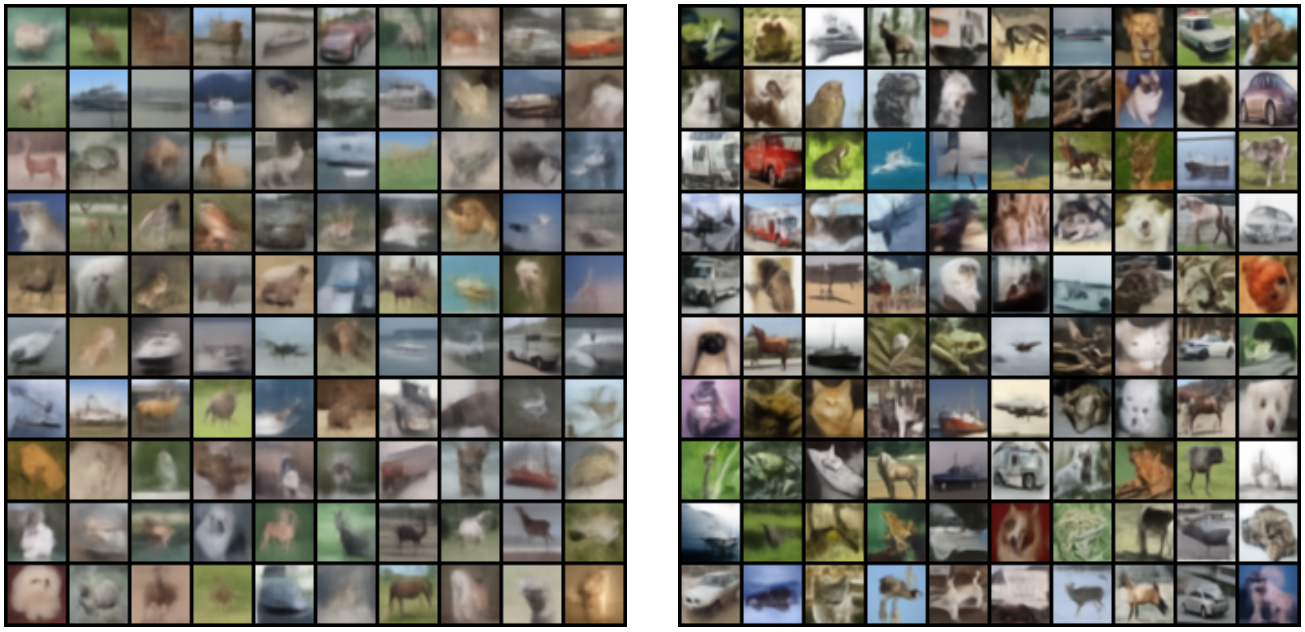

*Figure 6.* Unconditional OT samples: OT+FM (left) vs. OT+Iso-FM (right).