EBBFLOW: PINN-INSPIRED PROGRESSIVE CURRICULUM FOR MEANFLOW TRAINING

Anonymous authors
Paper under double-blind review

ABSTRACT

MeanFlow learns interval-average velocity for few-step generation, but its fixed temporal sampler mixes diagonal and interior supervision throughout training. We interpret temporal coordinates as collocation points: diagonal samples provide a stochastic boundary target, whereas finite intervals supply a differential-consistency target. The diagonal fraction controls how often the consistency component appears, while interval length enters its gradient through an explicit factor. Motivated by this distinction, EbbFlow uses Temporal Supervision Allocation (TSA) to schedule the diagonal fraction and interval geometry. It starts with diagonal supervision, then admits interior samples and expands their maximum interval; a left-anchored interval law preserves coverage across time. The per-sample loss, target, backbone, and inference procedure remain unchanged. In matched class-conditional ImageNet 256×256 runs without classifier-free guidance (DiT-B/2, one training seed per method, 400k updates), EbbFlow achieves 50k-sample FID of 39.65 at one network evaluation and 36.02 at two-evaluation ODE sampling, versus 44.82 and 39.86 for MeanFlow. Gradient probes show a mid-training difference in component geometry. Ablations support allocation scheduling and left-anchored intervals, although a fixed full-horizon variant performs better at one evaluation. These results suggest that temporal sampling is a useful training control whose benefits depend on the inference protocol.

1 INTRODUCTION

Generative flow models transport noise to data by learning a velocity field. MeanFlow (Geng et al., 2025a) replaces instantaneous velocity with an average velocity over a finite interval, making one- and two-step generation possible without changing the backbone. Its training coordinates nevertheless mix two qualitatively different cases: diagonal points $(r, t) = (t, t)$ recover ordinary flow-matching supervision, whereas interior points $r < t$ impose a self-referential finite-interval identity. The default ImageNet recipe fixes the diagonal fraction at 0.75 and exposes the model to the full interval range from the first update. These constants do not adapt as the model and its derivative-based target evolve.

We expose this choice by viewing temporal coordinates as collocation points on the triangle $0 \leq r \leq t \leq 1$. This reveals two independent controls. The diagonal fraction ρ determines *how many* samples receive only the stable boundary condition; the interval length $h = t - r$ determines *how strongly* each interior sample carries the differential component, through an explicit factor of h. Interval-binned gradient probes confirm that the two components are antagonistic and that their directional conflict and scale imbalance vary with both interval length and training progress. A fixed collocation measure therefore presents the same mixture of constraints even when their optimization geometry is changing.

Motivated by curriculum remedies for ill-conditioned physics-informed neural networks (PINNs; Krishnapriyan et al., 2021), we propose EbbFlow, which uses Temporal Supervision Allocation (TSA) to schedule its training coordinates. TSA first emphasizes the diagonal and short intervals, then progressively admits more interior samples and expands the interval horizon. A left-anchored construction preserves global coverage of time throughout the curriculum. Only the sampling measure changes: the MeanFlow per-sample loss, target construction, backbone, guidance, and inference procedure remain untouched. This also distinguishes TSA from AlphaFlow (Zhang et al.,

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

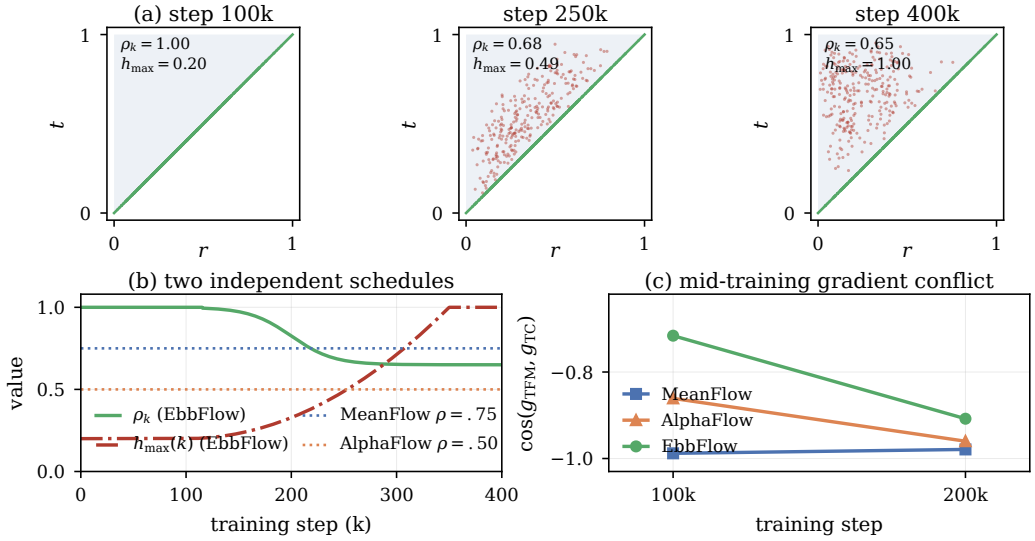


Figure 1: EbbFlow overview: (a) sampled temporal coordinates; (b) diagonal allocation ρ_k and horizon $h_{\max}(k)$; (c) mean component-gradient cosine at 100k and 200k. Figure 3 gives the full 100k–400k comparison.

2026), which changes the target construction; the two mechanisms are conceptually orthogonal. Short-to-long interval curricula also appear in gap-dependent weighting (Kim et al., 2026) and pre-trained speech enhancement (Wang et al., 2025). Our contribution is to separate and jointly schedule diagonal allocation and interval geometry within the sampling measure.

Our contributions are threefold. First, we formulate MeanFlow training as a boundary–interior collocation problem and separate allocation from interval geometry. Second, we connect the interval length to consistency-gradient scale and measure the resulting conflict throughout training. Third, in a matched single-seed ImageNet 256×256 comparison with DiT-B/2, EbbFlow improves FID from 44.82 to 39.65 at 1-NFE and from 39.86 to 36.02 at 2-NFE. Ablations show that both scheduling axes matter, while a fixed full-horizon left-anchored sampler remains better at 1-NFE. We accordingly present EbbFlow as a complementary optimization recipe, not a universally optimal schedule.

2 Related Work

Few-step generation and flow maps. Diffusion and flow-matching models learn transport from noise to data (Ho et al., 2020; Song et al., 2021; Lipman et al., 2023; Liu et al., 2023). Latent representations and transformer backbones make these models practical for image synthesis (Rombach et al., 2022; Peebles & Xie, 2023; Esser et al., 2024). Few-step generation reduces repeated network evaluations through teacher distillation (Salimans & Ho, 2022; Yin et al., 2024), trajectory consistency (Song et al., 2023; Song & Dhariwal, 2024; Geng et al., 2025b; Lu & Song, 2025), or learned two-time flow maps (Kim et al., 2024; Boffi et al., 2025; Frans et al., 2025; Zhou et al., 2025). EbbFlow instead retains MeanFlow's transport identity and changes its training coordinates.

MeanFlow and its variants. MeanFlow parameterizes transport by an average velocity whose differential identity reduces to flow matching at $r = t$ (Geng et al., 2025a). SplitMeanFlow replaces differentiation with interval-splitting consistency (Guo et al., 2025); improved MeanFlow changes the loss parameterization and guidance conditioning (Geng et al., 2026). Other variants modulate target gradients (You et al., 2025), decouple timestep conditioning (Lee et al., 2025), model ODE solutions (Luo et al., 2026), straighten couplings (Zhang et al., 2025), or reduce target variance (Lu & Wang, 2026). Discrete MeanFlow extends generation to finite state spaces through transition kernels (Khan et al., 2026). EbbFlow keeps the continuous MeanFlow target and parameterization fixed, scheduling where its identity is enforced.

Training dynamics and MeanFlow curricula. AlphaFlow identifies opposed trajectory-flow-matching and consistency gradients and anneals the target between these objectives (Zhang et al., 2026). Temporal Equilibrium MeanFlow (TEMF; Tu et al., 2026) addresses temporal scale imbalance through a fixed interval-dependent loss weight and a sigmoid schedule that increases the off-diagonal fraction, shifting from boundary stabilization to trajectory learning. Thus, boundary-to-interior allocation scheduling already appears in prior work. Kim et al. (2026) show that accurate instantaneous velocities support average-velocity learning: short gaps help their formation, whereas large gaps can disrupt it. Their method accelerates instantaneous-velocity learning and progressively relaxes gap-dependent loss weighting. Related curricula introduce trajectory-based mid-training (Hu et al., 2026), anneal a target discretization step (Hsu & Wood, 2026), or warm up gradient modulation (You et al., 2025). MeanSE initializes from a pretrained speech flow model and increases the maximum interval in stages from 0.2 to 1 (Wang et al., 2025). Short-to-long curricula are therefore also established. EbbFlow studies allocation together with a left-anchored interval law and a progressively expanded sampling horizon, retaining the baseline loss rather than introducing TEMF's interval weight. Its collocation view and interval-binned gradient diagnostics examine these sampling choices; the present experiments do not establish superiority over TEMF.

Timestep sampling, loss weighting, and curricula. Diffusion training distributes effort over noise levels through loss weights (Hang et al., 2023) or sampling densities, including the logit-normal law used in rectified-flow training (Esser et al., 2024). With matching density-ratio weights and sufficient proposal support, importance resampling preserves the population objective but changes estimator variance (Kingma & Gao, 2023). TSA instead changes the sampling measure without compensation; Section 4.4 discusses its relation to adaptive loss normalization. Related work treats diffusion timesteps as interacting tasks (Go et al., 2023), schedules denoising difficulty (Kim et al., 2025; Bengio et al., 2009), and anneals consistency discretization (Song & Dhariwal, 2024; Geng et al., 2025b). TSA schedules both allocation and interval geometry on MeanFlow's two-dimensional coordinate triangle.

PINN optimization and collocation. PINNs jointly fit boundary data and interior differential residuals (Raissi et al., 2019). Their gradients can be imbalanced (Wang et al., 2021), and differential constraints can produce ill-conditioned loss landscapes (Krishnapriyan et al., 2021; Rathore et al., 2024). Remedies include curriculum warm starts (Krishnapriyan et al., 2021), causality-respecting training (Wang et al., 2024; Mattey & Ghosh, 2022), and adaptive collocation (Wu et al., 2023). Physics Informed Distillation learns the solution of a teacher diffusion model's probability-flow ODE (Tee et al., 2024); EbbFlow instead applies the optimization analogy to training MeanFlow from scratch. We examine coordinate-dependent gradient interactions in Figure 2, using diagnostics also found in multi-task learning (Chen et al., 2018; Yu et al., 2020; Fifty et al., 2021).

3 Analyzing the Temporal Sampling of MeanFlow

MeanFlow is never trained on a single coordinate. Every update draws a batch of temporal pairs (r, t), and the ImageNet recipe fixes how they are drawn: 75% of the batch sits on the diagonal $r = t$, while the remaining pairs cover the full interval range $h = t - r \in (0, 1)$ from the first step onward (Geng et al., 2025a). Zhang et al. (2026) ask why so much of the budget is spent on the diagonal. We ask a complementary question: *why should either constant be fixed at all?*

3.1 Preliminaries

Flow matching. On the linear path $x_t = (1 - t)\, x + t\, \epsilon$ with data x and noise ϵ, Flow Matching (Lipman et al., 2023) regresses the conditional velocity $\epsilon - x$ with a network $v_\theta(x_t, t)$, whose minimizer is the marginal field $v(x_t, t) = \mathbb{E}[\epsilon - x \mid x_t]$. This target does not depend on θ; the MeanFlow target below does.

MeanFlow. MeanFlow (Geng et al., 2025a) models the average velocity over a finite interval $[r, t]$ with $t > r$ along a marginal-velocity characteristic,

$$u(x_t, r, t) \triangleq \frac{1}{t - r} \int_r^t v(x_\tau, \tau)\, d\tau, \tag{1}$$

with $u(x_t, t, t) = v(x_t, t)$ by continuous extension. A network $u_\theta(x_t, r, t)$ predicts this quantity directly. Direct integration is intractable during training, so MeanFlow differentiates Equation 1 with respect to t to obtain the MeanFlow identity

$$u(x_t, r, t) = v(x_t, t) - (t - r)\,\frac{d}{dt}u(x_t, r, t), \tag{2}$$

where the total derivative is computed with a Jacobian–vector product (JVP). One-step generation follows from $x_0 = x_1 - u_\theta(x_1, 0, 1)$, i.e. from the single coordinate $(r, t) = (0, 1)$.

Write $h = t - r$ for the interval length and $\widetilde{v}_t$ for the sampled training velocity ($\epsilon - x$ without guidance, or the velocity returned by the baseline's training-time guidance construction). Let $\widetilde{\mathcal{D}}_t = \partial_t + \widetilde{v}_t^\top \nabla_{x_t}$ be the derivative along the sampled characteristic, evaluated with a JVP. MeanFlow trains u_θ against the stop-gradient target

$$y_\theta(x_t, r, t) = \mathrm{sg}\big[\widetilde{v}_t - h\,\widetilde{\mathcal{D}}_t u_\theta(x_t, r, t)\big], \qquad \ell_{\mathrm{MF}} = \big\|u_\theta(x_t, r, t) - y_\theta(x_t, r, t)\big\|_2^2. \tag{3}$$

Two properties of Equation 3 are used repeatedly below. First, away from $r = t$ the target is *self-referential*: it is assembled from the network's own directional derivative, whose accuracy must be assessed separately from the value of u_θ at the sampled coordinate. Second, at $h = 0$ the correction vanishes and ℓ_{MF} reduces exactly to the flow-matching loss with the θ-independent target $\widetilde{v}_t$. The diagonal $r = t$ is thus the only place where MeanFlow supervision does not depend on current parameters. The corresponding marginal-field equation and the baseline's adaptive loss weighting are given in Appendix A.

3.2 MeanFlow training as temporal collocation

Our MeanFlow baseline (Geng et al., 2025a) assigns 75% of the batch to diagonal samples, drawing t from $\mathrm{LN}(-0.4, 1)$ and setting $r = t$. For the remaining samples, it independently draws a_1, a_2 from the same logit-normal law and sets $(r, t) = (\min(a_1, a_2), \max(a_1, a_2))$. Two constants are embedded in this procedure: the fraction of samples placed on the diagonal, fixed at 0.75, and the distribution of interval lengths induced by the order statistics of two logit-normal variables, which places mass on the full range $h \in (0, 1)$ from the first step onward. Neither constant is derived from Equation 2, and neither responds to the fact that u_θ—and therefore the target in Equation 3—changes substantially over training. To make these questions precise we adopt a collocation view. Temporal coordinates live on the triangle $\Omega = \{(r, t) : 0 \le r \le t \le 1\}$, whose hypotenuse $\Gamma_0 = \{(t, t)\}$ we call the diagonal. Training draws coordinates from a step-dependent measure P_k on Ω and averages Equation 3. We factor that measure as

$$P_k \;=\; \rho_k \left(p_t \otimes \delta_{\{r=t\}}\right) \;+\; (1 - \rho_k)\, q_k(r, t), \tag{4}$$

where p_t is the configured time sampler, $\rho_k \in [0, 1]$ is the diagonal fraction at step k, and q_k is a distribution supported on the interior $r < t$. Under this split the loss reads

$$\ell_{\mathrm{MF}}(\theta; x_t, r, t) = \begin{cases} \big\|u_\theta(x_t, t, t) - \widetilde{v}_t\big\|_2^2, & r = t \quad \text{(boundary condition)}, \\ \big\|u_\theta(x_t, r, t) - \mathrm{sg}\big[\widetilde{v}_t - h\,\widetilde{\mathcal{D}}_t u_\theta(x_t, r, t)\big]\big\|_2^2, & r < t \quad \text{(interior residual)}, \end{cases} \tag{5}$$

and the batch objective is a single mean of Equation 5 over B coordinates drawn from P_k. This is a renaming, not a new objective: substituting $\rho_k = 0.75$ and the order-statistic law q^{mm} (for min–max) recovers standard MeanFlow exactly. The fraction ρ_k determines sample counts, rather than weighting two separately averaged losses. The boundary condition on Γ_0 and the differential constraint on $r < t$ mirror the two collocation types of a PINN.

Where the PINN analogy stops. The split in Equation 5 is structurally PINN-like, but it is not a physical boundary-value problem. A deterministic PINN is given its boundary value, whereas the diagonal MeanFlow target $\widetilde{v}_t$ is a sample whose conditional expectation is the unknown marginal velocity. Consequently, even an expressive predictor faces conditional regression noise on the diagonal. Interior targets add a second source of error because their stopped target uses the current network derivative. The analogy therefore motivates how we organize and schedule the two collocation types; it does not imply that the diagonal can be solved exactly or that PINN convergence results transfer to generative training.

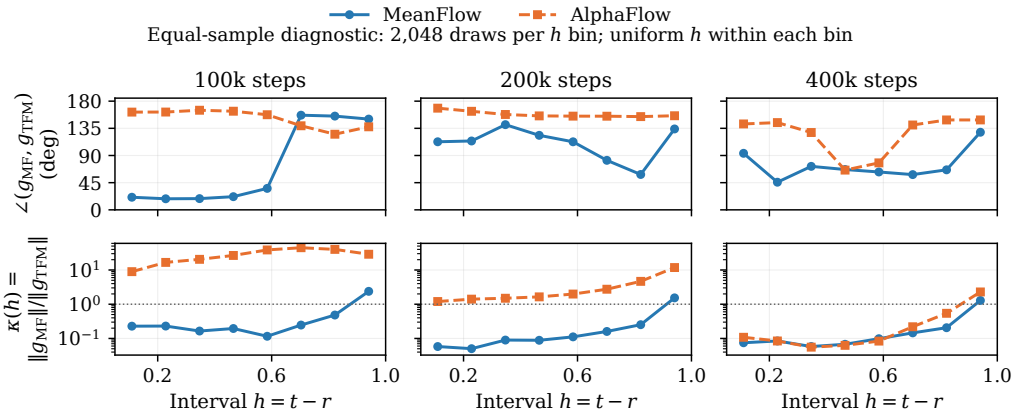


Figure 2: Full-gradient diagnostics for MeanFlow and AlphaFlow at 100k, 200k, and 400k. Top: $\angle(g_{\mathrm{MF}}, g_{\mathrm{TFM}})$ in degrees; bottom: $\kappa = \|g_{\mathrm{MF}}\|/\|g_{\mathrm{TFM}}\|$, unlike the component ratio q in Figure 3.

3.3 What the two constants control

The collocation view names the two constants but does not yet say what they act on. Let $a_\theta = u_\theta(x_t, r, t) - \widetilde{v}_t$ and $b_\theta = \mathrm{sg}[\widetilde{\mathcal{D}}_t u_\theta(x_t, r, t)]$, so that the interior branch is $\|a_\theta + h\, b_\theta\|_2^2$. Expanding,

$$\ell_{\mathrm{MF}} = \underbrace{\|a_\theta\|_2^2}_{\mathcal{L}_{\mathrm{TFM}}} + \underbrace{2h\,\langle a_\theta, b_\theta\rangle}_{\mathcal{L}_{\mathrm{TC}}} + \underbrace{h^2\|b_\theta\|_2^2}_{\text{no gradient}}, \qquad g_{\mathrm{TC}} = 2h\,(\nabla_\theta a_\theta)^\top b_\theta. \tag{6}$$

This is the trajectory-flow-matching / trajectory-consistency decomposition of Zhang et al. (2026), written per coordinate. The stopped-target gradient has an explicit factor of h:

$$\frac{\|g_{\mathrm{TC}}\|}{\|g_{\mathrm{TFM}}\|} = h \cdot \frac{\|(\nabla_\theta a_\theta)^\top b_\theta\|}{\|(\nabla_\theta a_\theta)^\top a_\theta\|} \tag{7}$$

The remaining factor also depends on the sampled coordinate, so this is not a monotonicity claim across intervals. The fraction ρ_k controls how many coordinates carry the consistency term; h enters its per-sample scale. The collocation view therefore shows that the two constants act on different quantities and can be varied independently; it does not prescribe whether either should vary during training.

3.4 The conflict is not stationary

Equation 6 fixes the structure of the conflict but not its magnitude, which depends on u_θ. We therefore probe it directly on raw DiT-B/2 checkpoints of MeanFlow and AlphaFlow at 100k, 200k and 400k steps (Figure 2). Standard MeanFlow keeps $\rho_k = 0.75$ throughout; the sigmoid AlphaFlow run keeps $\rho_k = 0.5$, where "sigmoid" refers to annealing its target parameter α from 1 to 0 rather than to its flow-matching ratio. For each of eight interval bins we measure the angle between $g_{\mathrm{MF}} = g_{\mathrm{TFM}} + g_{\mathrm{TC}}$ and g_{TFM}, and the norm ratio $\kappa(h) = \|g_{\mathrm{MF}}\|/\|g_{\mathrm{TFM}}\|$, using 2,048 samples in each of eight equal-width bins over $h \in [0.05, 1]$. For nonzero component and full gradients, let $c = \cos(g_{\mathrm{TFM}}, g_{\mathrm{TC}})$ and $q = \|g_{\mathrm{TC}}\|/\|g_{\mathrm{TFM}}\|$. Then $\kappa = \sqrt{1 + q^2 + 2qc}$ and $\cos\angle(g_{\mathrm{MF}}, g_{\mathrm{TFM}}) = (1 + qc)/\kappa$. Figure 2 reports the resulting full-gradient metrics; Figure 3 reports bin averages of c and q, rather than κ. Appendix A.3 gives the probe protocol. Both direction and scale vary with interval length and with training progress. At 100k, standard MeanFlow has a small full-gradient norm over most short and medium intervals, while AlphaFlow has a much larger norm and a wider angle; the longest intervals are the most sensitive in both methods. Because the probe uses equal per-bin budgets, it characterizes the coordinates themselves, not how often a sampler visits them. We return to the same probe in Section 5.4 to evaluate our schedules.

Design choices. These observations motivate three testable choices: (i) begin with the diagonal's θ-independent target before admitting interior samples (Section 4.1); (ii) control interval length

independently of the diagonal fraction (Section 4.2); and (iii) vary both controls to probe whether adaptation helps as gradient direction and scale change over training (Section 4.3). These choices affect only the sampling measure; the fixed-horizon control in Section 5.3 tests a limit of the joint schedule.

4 TEMPORAL SUPERVISION ALLOCATION

TSA replaces the fixed collocation measure with

$$P_k^{\mathrm{TSA}} = \rho_k \big(p_t \otimes \delta_{\{r=t\}}\big) + (1-\rho_k) q^{\mathrm{la}}_{h_{\max}(k)}, \tag{8}$$

where ρ_k is the diagonal fraction and q^{la}_H is a left-anchored interval law with horizon H. Relative to MeanFlow ($\rho_k = 0.75$, q^{mm}), TSA changes three sampling choices—allocation, interval parameterization, and horizon—whose effects Section 5.3 separates. The per-sample loss and target remain unchanged.

4.1 BOUNDARY-TO-INTERIOR ALLOCATION CURRICULUM

The diagonal provides the only θ-independent target (Section 3.1). Following the easy-to-hard optimization principle of Krishnapriyan et al. (2021), we first fit this anchor and progressively admit self-referential interior constraints. We use

$$\bar{\rho}_k = \rho_{\mathrm{start}} + (\rho_{\mathrm{end}} - \rho_{\mathrm{start}}) \sigma\left(\gamma \frac{k - k_{\mathrm{mid}}}{k_{\mathrm{end}} - k_{\mathrm{start}}}\right), \tag{9}$$

We implement $\rho_k = 1$ when $\bar{\rho}_k > 0.995$ and $\rho_k = \bar{\rho}_k$ otherwise, before per-device batch rounding. The reported run uses $(\rho_{\mathrm{start}}, \rho_{\mathrm{end}}) = (1, 0.65)$, $\gamma = 20$, and $(k_{\mathrm{start}}, k_{\mathrm{mid}}, k_{\mathrm{end}}) = (0, 200\mathrm{k}, 400\mathrm{k})$. Starting at $\rho = 1$ separates boundary fitting from interior learning; ending at 0.65 retains the anchor while increasing the interior fraction from MeanFlow's 25% to 35%.

For each per-device batch of size B, we place $\lfloor B\rho_k \rfloor$ samples on $r = t$, drawing t from the unchanged sampler p_t, and draw the remaining pairs from q_k; Appendix A.2 specifies the realized fractions and adaptive per-sample loss.

4.2 SHORT-TO-LONG INTERVAL CURRICULUM

Interval length controls the explicit factor $h = t - r$ in the consistency gradient (Equation 6). Capping h initially selects short-interval constraints that approach the diagonal as $h \to 0$, then admits longer intervals toward the inference coordinate $(0, 1)$.

Independent variables $\eta_h, \eta_r \sim \mathrm{LogitNormal}(-0.4, 1.0^2)$ define

$$h = h_{\max}(k)\eta_h, \qquad r = (1-h)\eta_r, \qquad t = r + h. \tag{10}$$

This guarantees $0 \le r \le t \le 1$ and $h \le h_{\max}(k)$. For any admissible h, r ranges over $[0, 1-h]$; collectively, the sampler retains support throughout $t \in [0, 1]$. It restricts interval length while preserving global time coverage.

The joint run uses the quadratic horizon

$$h_{\max}(k) = \begin{cases} 0.2, & k \le 100\mathrm{k}, \\ 0.2 + 0.8\left(\dfrac{k - 100\mathrm{k}}{250\mathrm{k}}\right)^2, & 100\mathrm{k} < k < 350\mathrm{k}, \\ 1, & k \ge 350\mathrm{k}. \end{cases} \tag{11}$$

Its convex shape delays large-h supervision relative to a linear ramp, motivated by the component-scale imbalance in Equation 6. This is a design choice: Section 5.3 compares it with a full horizon, but does not establish that quadratic growth is optimal.

Interval law versus horizon. Even at $h_{\max} = 1$, left anchoring differs from MeanFlow's min–max sampler. With 2×10^6 Monte Carlo draws per law, the mean interior interval is about 0.42 versus 0.23, and the mass near the inference corner ($r < 0.1$, $t > 0.9$) is 1.8% versus 0.035%. Greater exposure near deployment may explain part of the parameterization benefit, although the mechanism is not isolated. To separate it from horizon scheduling, the full-horizon control uses the same ρ_k schedule and left-anchored law as TSA, with $h_{\max} \equiv 1$.

4.3 Joint schedule

Figure 1 shows the joint curriculum, which moves through boundary fitting ($k \leq 100$k), interior admission ($100\text{k} < k < 350\text{k}$), and full-horizon training ($k \geq 350$k). We expect reduced gradient conflict during interior admission. Once the schedules settle, TSA and the baselines face the same constraint family with different allocation and interval laws, so the advantage need not persist (Section 5.4). A coordinate-level exposure proxy, $S_k = (1 - \rho_k)\mathbb{E}_{q_k}[t - r]$, shows that TSA supplies $0.41\times$ MeanFlow's interior exposure at 200k and $2.51\times$ at 400k, while its cumulative exposure through 400k is 88.2% of MeanFlow's (Appendix B). This proxy measures coordinate exposure rather than loss or gradient strength; the schedule changes its timing, distribution, and cumulative amount.

4.4 Design choices, cost, and scope

Global coverage. Restricting t may induce covariate shift as the shared generative network moves between marginals; this explanation remains a hypothesis. Left anchoring instead restricts h while preserving global time coverage.

Resampling and target construction. For a fixed per-sample loss, changing the sampling density can be represented by density-ratio loss weights in expectation, provided the proposal covers the required support (Kingma & Gao, 2023). This does not imply identical estimator variance or identical behavior when weights are inserted before adaptive loss normalization (Kim et al., 2026). TSA changes the training measure without importance compensation, retaining the baseline's single adaptively weighted batch mean (Appendix A.2). It schedules sample counts and interval support while preserving the per-sample loss; we do not claim resampling is universally better than weighting. AlphaFlow instead anneals the target between finite-difference and continuous derivatives (Zhang et al., 2026). We fix $\alpha = 0$; combining the two interventions is possible but untested, and their gains need not be additive.

Training step. Algorithm 1 adds two schedules and draws pairs (r, t). It adds no model parameters or inference steps. We still compute JVPs on the full batch, so we claim no faster training. Here LN is the baseline logit-normal law; $\widetilde{v}_i$ is the conditional velocity, and ℓ_{base} is the baseline loss.

Scope. TSA changes the training path but leaves gradient conflict in place. We test this with different interval laws and inference samplers; the horizon rule does not improve every protocol.

Algorithm 1: EbbFlow: TSA training step.

$H \leftarrow h_{\max}(k), \quad n \leftarrow \lfloor B\rho_k \rfloor, \quad L \leftarrow 0$
for $i = 1, \ldots, B$:
  if $i \leq n$: $t \sim p_t, \quad r \leftarrow t$
  else:
    $a, b \overset{\text{iid}}{\sim} \text{LN}(-0.4, 1)$
    $h \leftarrow Ha, \quad r \leftarrow (1 - h)b, \quad t \leftarrow r + h$
  $x_t \leftarrow (1 - t)x_i + t\epsilon_i$
  $\widetilde{v}_i \leftarrow \text{baseline_velocity}(x_i, \epsilon_i, t)$
  $u \leftarrow u_\theta(x_t, r, t)$
  $g \leftarrow \text{JVP}(u_\theta; (x_t, r, t), (\widetilde{v}_i, 0, 1))$
  $y \leftarrow \text{stopgrad}(\widetilde{v}_i - (t - r)g)$
  $L \leftarrow L + \ell_{\text{base}}(u, y)/B$
$\text{update}(\theta, L)$

5 EXPERIMENTS

5.1 EVALUATION PROTOCOL

All matched quality runs use class-conditional latent ImageNet (Deng et al., 2009) 256×256 generation with DiT-B/2 (131M parameters; Peebles & Xie, 2023), a global batch of 256, and 400k updates. MeanFlow keeps $\rho_k = 0.75$ and uses min–max intervals q^{mm}. AlphaFlow uses sigmoid target annealing with $\rho_k = 0.5$. EbbFlow retains the MeanFlow target ($\alpha = 0$) and uses Equation 8. The runs share the optimizer, EMA rule, checkpoint selection, random class sampling, and evaluation code. All matched runs are class-conditional but use no classifier-free guidance (Ho & Salimans, 2022): the training target uses the conditional velocity, and the inference guidance scale is one. We report FID (Heusel et al., 2017) and DINO-FID (FDD) from 50k generated images at one and two network evaluations (NFE). FDD uses DINOv2-ViT-L/14 features (Stein et al., 2023). The two-NFE results in Table 1 use ODE sampling.

Each matched training configuration has one seed. The AlphaFlow row averages two independent 50k-sample evaluations of the same checkpoint; it is not a two-seed training average. Published results with different backbones, budgets, or protocols are not directly comparable.

Table 1: Matched ImageNet 256×256 results (DiT-B/2, 400k updates; 50k-sample FID). Two-NFE uses ODE sampling. Full metrics: Table 4.

Method	1-NFE	2-NFE ODE
MeanFlow	44.82	39.86
AlphaFlow	41.89	37.67
EbbFlow	**39.65**	**36.02**

Table 2: DiT-B/2 ablations at 400k updates. MM: min–max; LA: left anchor; Q: quadratic horizon (otherwise full). Full metrics: Table 5.

	ρ_k	Sampler	1-NFE	2-NFE ODE
(a)	0.75	MM	44.82	39.86
(b)	0.82525	MM	44.13	38.77
(c)	$1 \to 0.65$	MM	41.80	37.22
(d)	0.75	LA-Q	40.78	36.92
(e)	$1 \to 0.65$	LA	**38.85**	36.25
(f)	$1 \to 0.65$	LA-Q	39.65	**36.02**

5.2 MATCHED IMAGENET RESULTS

Under this protocol, EbbFlow lowers FID relative to matched MeanFlow by 5.17 points (11.5%) at one NFE and 3.84 points (9.6%) at two NFE. Relative to AlphaFlow, the reductions are 2.24 and 1.65 points. EbbFlow has the lowest two-NFE FDD, but its one-NFE FDD is 9.1 points *higher* than AlphaFlow's (Table 4). The metrics therefore disagree on the one-NFE ordering of those two methods. These single-seed comparisons do not estimate training-seed variability.

5.3 ATTRIBUTION OF TEMPORAL-COORDINATE CHOICES

Table 2 separates allocation, interval law, and horizon. The full-horizon left-anchor control shares EbbFlow's allocation schedule and interval law, isolating horizon expansion.

Allocation. With min–max intervals fixed, scheduling the allocation improves one-NFE FID by 3.02 points over MeanFlow (a→c) and by 2.33 points over the static-mixture control (b→c). With the quadratic left-anchor horizon fixed, the allocation schedule improves one-NFE FID by 1.13 points (d→f). The static control suggests that temporal allocation matters beyond an approximately matched average mixture. Per-device rounding gives slightly different realized diagonal fractions (about 82.03% static versus 82.29% scheduled), so this is not an exact discrete-budget control.

Interval law and horizon. At the same scheduled allocation, replacing min–max intervals (c) with full-horizon left-anchored intervals (e) lowers one-NFE FID by 2.95 points and two-NFE ODE FID by 0.97 points. This isolates the interval-law effect under scheduled allocation, not under fixed $\rho_k = 0.75$. Adding the quadratic horizon (e→f) *raises* one-NFE FID by 0.80 points and consistency

FID by 0.60 points (Table 5), but *lowers* two-NFE ODE FID by 0.23 points. Thus the horizon schedule is not uniformly beneficial. A linear-horizon control and multiple seeds would be needed to test whether its shape or the small two-NFE difference is robust.

Sampler dependence. With each checkpoint held fixed, two-NFE ODE sampling beats consistency sampling for every $\alpha = 0$ variant, by 3.12 to 6.01 FID (Table 5). Our two-NFE claims are therefore specific to ODE sampling and do not establish which sampler suits AlphaFlow's target.

5.4 Optimization-path comparison

We apply the equal-bin probe of Section 3.4 to raw 100k, 200k, 300k, and 400k checkpoints of all three methods (FP32, evaluation mode), using 2,048 samples per bin. It diagnoses coordinate-level gradient geometry rather than the training sampler's visit frequency.

At 100k, EbbFlow is effectively in its pure boundary-fitting phase. The probe nevertheless includes untrained interior coordinates, yielding a less-negative mean component cosine (-0.716) and a large component norm ratio q (48.9). This is task separation, not balanced joint optimization. At 200k, the mean q falls to 1.25, near MeanFlow's 1.18 and below AlphaFlow's 4.31. EbbFlow's mean cosine is also less negative (-0.908 versus -0.979 and -0.960). A smaller directional advantage persists at 300k. By 400k, all three mean cosines are about -0.98, and the mean norm ratios lie between 1.10 and 1.38. These measurements suggest a mid-training difference followed by similar bin-averaged diagnostics at 400k. This post-hoc probe cannot by itself establish that the path change caused the final FID differences.

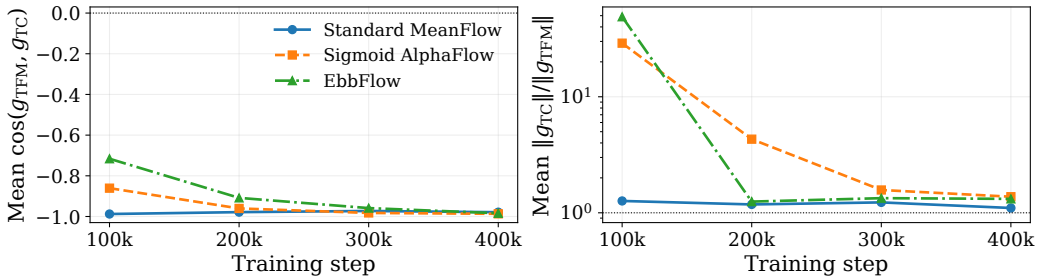


Figure 3: Component-gradient diagnostics at 100k–400k. Left: mean $\cos(g_{\mathrm{TFM}}, g_{\mathrm{TC}})$; right: mean $q = \|g_{\mathrm{TC}}\|/\|g_{\mathrm{TFM}}\|$. Each point averages eight interval bins.

6 Discussion and Conclusion

We presented EbbFlow, which schedules diagonal allocation and interval geometry in MeanFlow's temporal sampler. In matched single-seed DiT-B/2 runs at 400k updates, it improves MeanFlow's FID from 44.82 to 39.65 at 1-NFE and from 39.86 to 36.02 at 2-NFE ODE. At 200k, gradient probes suggest less component-gradient opposition than both baselines and better scale balance than AlphaFlow; by 400k, bin-averaged diagnostics are similar across methods. Ablations support benefits from allocation scheduling and left-anchored intervals, while the full-horizon left-anchor control outperforms the joint curriculum for two of three samplers. The evidence supports temporal sampling as a complementary training control and a 2-NFE benefit from the joint curriculum. Multi-seed matched runs remain necessary to establish reliability.

AI use statement

Generative AI tools assisted with finding and checking related literature, checking and clarifying parts of the mathematical analysis, reviewing the experimental methodology and interpretation of results, drafting and editing manuscript text, preparing code and figures, and checking formatting. The authors reviewed the AI-assisted analysis, text, code, and figures and take responsibility for the final content, claims, and artifacts in this submission.

Reproducibility statement

Section 4 and Algorithm 1 specify TSA; Appendices A.2–A.3, B, C, and D describe implementation details, the diagnostic protocol, schedules, complete evaluation metrics, and experimental configurations. Section 5 and Appendix D state the training-seed count, evaluation sample counts, checkpoint selection, sampling modes, batch sizes, and guidance settings.

References

Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pp. 41–48, 2009. URL https://doi.org/10.1145/1553374.1553380.

Nicholas M. Boffi, Michael S. Albergo, and Eric Vanden-Eijnden. Flow map matching with stochastic interpolants: A mathematical framework for consistency models. *Transactions on Machine Learning Research*, 2025. URL https://openreview.net/forum?id=cqDH0e6ak2.

Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. GradNorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International Conference on Machine Learning*, volume 80, pp. 794–803, 2018. URL https://proceedings.mlr.press/v80/chen18a.html.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009. URL https://doi.org/10.1109/CVPR.2009.5206848.

Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis. In *International Conference on Machine Learning*, volume 235, pp. 12606–12633, 2024. URL https://proceedings.mlr.press/v235/esser24a.html.

Christopher Fifty, Ehsan Amid, Zhe Zhao, Tianhe Yu, Rohan Anil, and Chelsea Finn. Efficiently identifying task groupings for multi-task learning. In *Advances in Neural Information Processing Systems*, volume 34, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/e77910ebb93b511588557806310f78f1-Abstract.html.

Kevin Frans, Danijar Hafner, Sergey Levine, and Pieter Abbeel. One step diffusion via shortcut models. In *International Conference on Learning Representations*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/559a0998fab1d19b80e7e43a5852401c-Abstract-Conference.html.

Zhengyang Geng, Mingyang Deng, Xingjian Bai, J. Zico Kolter, and Kaiming He. Mean flows for one-step generative modeling. In *Advances in Neural Information Processing Systems*, volume 38, 2025a. URL https://papers.neurips.cc/paper_files/paper/2025/hash/6d13e085b79d454da5910e4ca82a3d9d-Abstract-Conference.html.

Zhengyang Geng, Ashwini Pokle, Weijian Luo, Justin Lin, and J. Zico Kolter. Consistency models made easy. In *International Conference on Learning Representations*, 2025b. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/bb166dd4de5dba363bf1023eb956a826-Abstract-Conference.html.

Zhengyang Geng, Yiyang Lu, Zongze Wu, Eli Shechtman, J. Zico Kolter, and Kaiming He. Improved mean flows: On the challenges of fastforward generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 30467–30476, 2026. URL https://openaccess.thecvf.com/content/CVPR2026/html/Geng_Improved_Mean_Flows_On_the_Challenges_of_Fastforward_Generative_Models_CVPR_2026_paper.html.

Hyojun Go, JinYoung Kim, Yunsung Lee, Seunghyun Lee, Shinhyeok Oh, Hyeongdon Moon, and Seungtaek Choi. Addressing negative transfer in diffusion models. In *Advances in Neural Information Processing Systems*, volume 36, 2023. URL https://papers.neurips.cc/paper_files/paper/2023/hash/56a7b9a07ae01ddea762dcc51280298b-Abstract-Conference.html.

Yi Guo, Wei Wang, Zhihang Yuan, Rong Cao, Kuan Chen, Zhengyang Chen, Yuanyuan Huo, Yang Zhang, Yuping Wang, Shouda Liu, and Yuxuan Wang. SplitMeanFlow: Interval splitting consistency in few-step generative modeling. *arXiv preprint arXiv:2507.16884*, 2025. URL https://arxiv.org/abs/2507.16884.

Tiankai Hang, Shuyang Gu, Chen Li, Jianmin Bao, Dong Chen, Han Hu, Xin Geng, and Baining Guo. Efficient diffusion training via min-SNR weighting strategy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 7441–7451, 2023. URL https://openaccess.thecvf.com/content/ICCV2023/html/Hang_Efficient_Diffusion_Training_via_Min-SNR_Weighting_Strategy_ICCV_2023_paper.html.

Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://papers.nips.cc/paper/2017/hash/8a1d694707eb0fefe65871369074926d-Abstract.html.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. URL https://arxiv.org/abs/2207.12598.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html.

Chia-Hong Hsu and Frank Wood. Discrete Meanflow training curriculum. *arXiv preprint arXiv:2604.08837*, 2026. URL https://arxiv.org/abs/2604.08837.

Zheyuan Hu, Chieh-Hsin Lai, Yuki Mitsufuji, and Stefano Ermon. CMT: Mid-training for efficient learning of consistency, mean flow, and flow map models. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/fb44a668c2d4bc984e9d6ca261262cbb-Abstract-Conference.html.

Fairoz Nower Khan, Nabuat Zaman Nahim, Md Sajid Ahmed, Ruiquan Huang, and Peizhong Ju. Discrete MeanFlow: One-step generation via conditional transition kernels. *arXiv preprint arXiv:2605.12805*, 2026. URL https://arxiv.org/abs/2605.12805.

Dongjun Kim, Chieh-Hsin Lai, Wei-Hsiang Liao, Naoki Murata, Yuhta Takida, Toshimitsu Uesaka, Yutong He, Yuki Mitsufuji, and Stefano Ermon. Consistency trajectory models: Learning probability flow ODE trajectory of diffusion. In *International Conference on Learning Representations*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/c204d12afa0175285e5aac65188808b4-Abstract-Conference.html.

Jin-Young Kim, Hyojun Go, Soonwoo Kwon, and Hyun-Gyoon Kim. Denoising task difficulty-based curriculum for training diffusion models. In *International Conference on Learning Representations*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/a1cb36e60c6c6f0b1c66a042580040d4-Abstract-Conference.html.

Jin-Young Kim, Hyojun Go, Lea Bogensperger, Julius Erbach, Nikolai Kalischek, Federico Tombari, Konrad Schindler, and Dominik Narnhofer. Understanding, accelerating, and improving MeanFlow training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 37992–38003, 2026. URL https://openaccess.thecvf.

com/content/CVPR2026/html/Kim_Understanding_Accelerating_and_Improving_MeanFlow_Training_CVPR_2026_paper.html.

Diederik P. Kingma and Ruiqi Gao. Understanding diffusion objectives as the ELBO with simple data augmentation. In *Advances in Neural Information Processing Systems*, volume 36, 2023. URL https://papers.neurips.cc/paper_files/paper/2023/hash/ce79fbf9baef726645bc2337abb0ade2-Abstract-Conference.html.

Aditi S. Krishnapriyan, Amir Gholami, Shandian Zhe, Robert M. Kirby, and Michael W. Mahoney. Characterizing possible failure modes in physics-informed neural networks. In *Advances in Neural Information Processing Systems*, volume 34, 2021. URL https://papers.neurips.cc/paper_files/paper/2021/hash/df438e5206f31600e6ae4af72f2725f1-Abstract.html.

Kyungmin Lee, Sihyun Yu, and Jinwoo Shin. Decoupled MeanFlow: Turning flow models into flow maps for accelerated sampling. *arXiv preprint arXiv:2510.24474*, 2025. URL https://arxiv.org/abs/2510.24474.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=PqvMRDCJT9t.

Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=XVjTT1nw5z.

Cheng Lu and Yang Song. Simplifying, stabilizing and scaling continuous-time consistency models. In *International Conference on Learning Representations*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/7e9c2053258b1bdd32ff2654802cd594-Abstract-Conference.html.

Juanwu Lu and Ziran Wang. On variance reduction in learning mean flows. *arXiv preprint arXiv:2605.09235*, 2026. URL https://arxiv.org/abs/2605.09235.

Tianze Luo, Haotian Yuan, and Zhuang Liu. SoFlow: Solution flow models for one-step generative modeling. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/a654298de44ec95dec28aa37e8d0af51-Abstract-Conference.html.

Revanth Mattey and Susanta Ghosh. A novel sequential method to train physics informed neural networks for Allen Cahn and Cahn Hilliard equations. *Computer Methods in Applied Mechanics and Engineering*, 390:114474, 2022. URL https://doi.org/10.1016/j.cma.2021.114474.

William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4195–4205, 2023. URL https://openaccess.thecvf.com/content/ICCV2023/html/Peebles_Scalable_Diffusion_Models_with_Transformers_ICCV_2023_paper.html.

Maziar Raissi, Paris Perdikaris, and George E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. URL https://doi.org/10.1016/j.jcp.2018.10.045.

Pratik Rathore, Weimu Lei, Zachary Frangella, Lu Lu, and Madeleine Udell. Challenges in training PINNs: A loss landscape perspective. In *International Conference on Machine Learning*, volume 235, pp. 42159–42191, 2024. URL https://proceedings.mlr.press/v235/rathore24a.html.

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. URL https://openaccess.thecvf.com/content/CVPR2022/html/Rombach_

High-Resolution_Image_Synthesis_With_Latent_Diffusion_Models_CVPR_2022_paper.html.

Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=TIdIXIpzhoI.

Yang Song and Prafulla Dhariwal. Improved techniques for training consistency models. In *International Conference on Learning Representations*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/41bd71e7bf7f9fe68f1c936940fd06bd-Abstract-Conference.html.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=PxTIG12RRHS.

Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *International Conference on Machine Learning*, volume 202, pp. 32211–32252, 2023. URL https://proceedings.mlr.press/v202/song23a.html.

George Stein, Jesse C. Cresswell, Rasa Hosseinzadeh, Yi Sui, Brendan Leigh Ross, Valentin Villecroze, Zhaoyan Liu, Anthony L. Caterini, J. Eric T. Taylor, and Gabriel Loaiza-Ganem. Exposing flaws of generative model evaluation metrics and their unfair treatment of diffusion models. In *Advances in Neural Information Processing Systems*, volume 36, 2023. URL https://papers.neurips.cc/paper_files/paper/2023/hash/0bc795afae289ed465a65a3b4b1f4eb7-Abstract-Conference.html.

Joshua Tian Jin Tee, Kang Zhang, Hee Suk Yoon, Dhananjaya Nagaraja Gowda, Chanwoo Kim, and Chang D. Yoo. Physics informed distillation for diffusion models. *Transactions on Machine Learning Research*, 2024. URL https://openreview.net/forum?id=rOvaUsF996.

Yuanpeng Tu, Yunpeng Chen, Xinyu Zhang, Chao Liao, and Hengshuang Zhao. Temporal equilibrium MeanFlow: Bridging the scale gap for one-step generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16064–16073, 2026. URL https://openaccess.thecvf.com/content/CVPR2026/html/Tu_Temporal_Equilibrium_MeanFlow_Bridging_the_Scale_Gap_for_One-Step_Generation_CVPR_2026_paper.html.

Jiahe Wang, Hongyu Wang, Wei Wang, Lei Yang, Chenda Li, Wangyou Zhang, Lufen Tan, and Yanmin Qian. MeanSE: Efficient generative speech enhancement with mean flows. *arXiv preprint arXiv:2509.21214*, 2025. URL https://arxiv.org/abs/2509.21214.

Sifan Wang, Yujun Teng, and Paris Perdikaris. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, 43(5): A3055–A3081, 2021. URL https://doi.org/10.1137/20M1318043.

Sifan Wang, Shyam Sankaran, and Paris Perdikaris. Respecting causality for training physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 421:116813, 2024. URL https://doi.org/10.1016/j.cma.2024.116813.

Chenxi Wu, Min Zhu, Qinyang Tan, Yadhu Kartha, and Lu Lu. A comprehensive study of non-adaptive and residual-based adaptive sampling for physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 403:115671, 2023. URL https://doi.org/10.1016/j.cma.2022.115671.

Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Frédo Durand, William T. Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6613–6623, 2024. URL https://openaccess.thecvf.com/content/CVPR2024/html/Yin_One-step_Diffusion_with_Distribution_Matching_Distillation_CVPR_2024_paper.html.

Haochen You, Baojing Liu, and Hongyang He. Modular MeanFlow: Towards stable and scalable one-step generative modeling. *arXiv preprint arXiv:2508.17426*, 2025. URL https://arxiv.org/abs/2508.17426.

Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. In *Advances in Neural Information Processing Systems*, volume 33, 2020. URL https://papers.neurips.cc/paper_files/paper/2020/hash/3fe78a8acf5fda99de95303940a2420c-Abstract.html.

Huijie Zhang, Aliaksandr Siarohin, Willi Menapace, Michael Vasilkovsky, Sergey Tulyakov, Qing Qu, and Ivan Skorokhodov. AlphaFlow: Understanding and improving MeanFlow models. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/e8c20cafe841cba3e31a17488dc9c3f1-Abstract-Conference.html.

Xinxi Zhang, Shiwei Tan, Quang Nguyen, Quan Dao, Ligong Han, Xiaoxiao He, Tunyu Zhang, Chengzhi Mao, Dimitris Metaxas, and Vladimir Pavlovic. Overcoming the curvature bottleneck in MeanFlow. *arXiv preprint arXiv:2511.23342*, 2025. URL https://arxiv.org/abs/2511.23342.

Linqi Zhou, Stefano Ermon, and Jiaming Song. Inductive moment matching. In *International Conference on Machine Learning*, volume 267, pp. 78651–78686, 2025. URL https://proceedings.mlr.press/v267/zhou25c.html.

A Detailed PINN-like formulation

A.1 Characteristic constraint and diagonal trace

Let x_0 denote a data sample, ϵ a noise sample, and $x_t = (1-t)x_0 + t\epsilon$. A model $u_\theta(x_t, r, t)$ predicts transport over an interval $0 \le r \le t \le 1$. Holding r fixed, define the characteristic derivative

$$\mathcal{D}_t \triangleq \partial_t + v(x_t, t)^\top \nabla_{x_t}. \tag{12}$$

The ideal MeanFlow identity in Equation 2 motivates the interior constraint

$$\mathcal{R}_\theta(x_t, r, t) \triangleq u_\theta(x_t, r, t) - v(x_t, t) + (t-r)\mathcal{D}_t u_\theta(x_t, r, t) = 0, \qquad r < t. \tag{13}$$

On the diagonal boundary $\Gamma_0 = \{r = t\}$, the coefficient of the derivative term vanishes, yielding

$$\mathcal{B}_\theta(x_t, t) \triangleq u_\theta(x_t, t, t) - v(x_t, t) = 0. \tag{14}$$

The same structure has an initial-value representation. Define the displacement $S_\theta(x_t, r, t) = (t-r)u_\theta(x_t, r, t)$. By the product rule, Equation 13 becomes

$$\mathcal{D}_t S_\theta(x_t, r, t) = v(x_t, t), \qquad S_\theta(x_r, r, r) = 0. \tag{15}$$

The zero-displacement trace holds for any bounded u_θ; by itself it does not enforce Equation 14. Together with the interior constraint and sufficient regularity, it recovers the diagonal relation.

This is a PINN-like optimization view rather than a conventional deterministic PINN. The states x_t are sampled from data–noise interpolations, and the unknown marginal velocity $v(x_t, t)$ is supervised using conditional velocities. The view is nevertheless useful because it separates diagonal boundary collocation from interior differential-constraint collocation, which are precisely the two types of samples controlled by our schedules.

A.2 Stopped target and implementation correspondence

We retain the continuous MeanFlow target and modify only temporal sampling. In the AlphaFlow implementation, this corresponds to fixing $\alpha = 0$, so the finite-difference branch is inactive. Let $\widetilde{v}_t$ denote the training velocity: $\widetilde{v}_t = \epsilon - x_0$ without guidance, or the velocity returned by the baseline's training-time guidance construction. For each sample, the implementation computes

$$g_\theta(x_t, r, t) = \mathrm{JVP}\big(u_\theta, (x_t, r, t), (\widetilde{v}_t, 0, 1)\big), \tag{16}$$

$$y_\theta(x_t, r, t) = \mathrm{sg}\big[\widetilde{v}_t - (t-r)g_\theta(x_t, r, t)\big], \tag{17}$$

where sg stops gradients through the target. When $r = t$, the JVP contribution vanishes and $y_\theta = \widetilde{v}_t$. For a d-dimensional output, the baseline's adaptively weighted per-sample loss is

$$e_i = \frac{1}{d} \left\| u_\theta(x_{t_i}, r_i, t_i) - y_\theta(x_{t_i}, r_i, t_i) \right\|_2^2, \qquad \ell_i = \frac{e_i}{\mathrm{sg}(e_i) + \varepsilon_{\mathrm{w}}}, \quad \varepsilon_{\mathrm{w}} = 10^{-3}. \tag{18}$$

The conditional tangent and stopped target form a stochastic training surrogate for the ideal marginal-field constraint. Per sample, the conditional velocity need not equal the marginal velocity; stopping the JVP target also gives a different parameter derivative from direct minimization of $\|\mathcal{R}_\theta\|^2$.

For a batch of size B, the code assigns

$$B_{\mathrm{FM}} = \lfloor B\rho_k \rfloor, \qquad B_{\mathrm{int}} = B - B_{\mathrm{FM}}, \qquad \widehat{\rho}_k = B_{\mathrm{FM}}/B, \tag{19}$$

and averages ℓ_i over all samples. Hence ρ_k changes sample counts, not an external coefficient applied after separately averaging two losses. Diagonal times follow the fixed configured logit-normal sampler, while off-diagonal pairs follow $q_k(r, t)$. Table 3 records the notation used by the implementation.

Table 3: Correspondence between the paper notation and the AlphaFlow training implementation.

Paper notation	Implementation quantity
r	`t_next` or `sigma_next`
t	`t` or `sigma`
x_t	`x_t`
$\widetilde{v}_t$	`velocity_cfg`
$u_\theta(x_t, r, t)$	`pred_mean_velocity`
g_θ	`videos_dudt_mf`
y_θ	output of `_compute_mean_velocity_c`
ρ_k	`ratio_fm`

A.3 Gradient-probe protocol

Both diagnostic figures use the continuous MeanFlow target ($\alpha = 0$), including for AlphaFlow checkpoints; they do not measure AlphaFlow's annealed training target. Raw checkpoints run in FP32 and evaluation mode. Eight equal-width bins partition $h \in [0.05, 1]$; each uses 2,048 samples, with h uniform within the bin and r uniform on $[0, 1 - h]$. The seed is 20260907. Sample indices, noise, and class conditions match across methods, but latent banks are generated separately and need not be identical. Each method reuses its bank across bins and checkpoints. Gradients of the two components include the common stopped adaptive weight from Equation 18 and are averaged over samples before computing direction and norm statistics. No sampling-frequency weights enter. Figure 2 transforms these statistics to full-gradient metrics in each bin; Figure 3 takes arithmetic means of the component metrics over bins. Curves are point estimates, without uncertainty estimates across training seeds.

B TSA schedule details

Schedule values. At 100k, 200k, 300k, and 400k updates, the clamped diagonal fractions before per-device batch rounding are 1, 0.825, 0.652, and 0.650, and the horizons are 0.2, 0.328, 0.712, and 1.0. During boundary fitting ($k \leq 100$k) the horizon is inactive, so a probe of the still-untrained interior coordinates can have a large norm ratio.

Effective supervision. We summarize the coordinate-level exposure to the consistency component by

$$S_k = (1 - \rho_k)\mathbb{E}_{(r,t)\sim q_k}[t - r]. \tag{20}$$

This pre-rounding proxy combines interior frequency with mean interval length; it is neither a measured loss nor a gradient norm. TSA supplies $0.41\times$ MeanFlow's instantaneous exposure at 200k

and $2.51\times$ at 400k, with cumulative exposure of 88.2% through 400k (Figure 4). The late ratio combines $1.4\times$ more interior samples with an approximately $1.8\times$ longer mean interval.

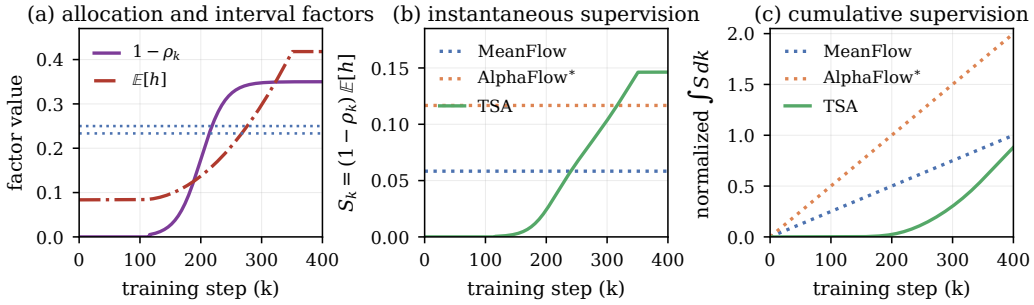


Figure 4: Sampler-based exposure before batch rounding: (a) interior fraction and mean interval; (b) their product S_k; (c) cumulative exposure relative to MeanFlow's 400k total. AlphaFlow* denotes coordinate exposure, not gradient strength.

C Complete generation metrics

Tables 4 and 5 retain the full metrics corresponding to the compact main-text tables. Each configuration has one training seed; all evaluations use 50k generated samples.

Table 4: Complete matched ImageNet 256×256 results with DiT-B/2 at 400k updates. AlphaFlow averages two evaluations of the same checkpoint. Lower FID and FDD are better; two-NFE results use ODE sampling.

	1-NFE		2-NFE ODE	
Method	FID	FDD	FID	FDD
MeanFlow	44.82	843.8	39.86	805.7
AlphaFlow	41.89	**787.8**	37.67	772.2
EbbFlow	**39.65**	796.9	**36.02**	**768.2**

Table 5: Complete DiT-B/2 ablations at 400k updates, measured by 50k-sample FID. MM denotes min–max intervals and LA denotes left anchoring. Static $\rho = 0.82525$ approximates the time-averaged allocation of $1 \to 0.65$. Cons. denotes two-NFE consistency sampling.

	ρ_k	Law	Horizon	1-NFE	2-NFE ODE	2-NFE cons.
(a)	0.75	MM	full	44.82	39.86	45.51
(b)	0.82525	MM	full	44.13	38.77	44.78
(c)	$1 \to 0.65$	MM	full	41.80	37.22	42.35
(d)	0.75	LA	quadratic	40.78	36.92	41.67
(e)	$1 \to 0.65$	LA	full	**38.85**	36.25	**39.37**
(f)	$1 \to 0.65$	LA	quadratic	39.65	**36.02**	39.97

D DETAILED EXPERIMENTAL CONFIGURATIONS

Table 6 lists the shared architecture, training, and evaluation settings of the matched DiT-B/2 experiments. Table 7 specifies the temporal samplers. These settings describe the matched DiT-B/2 runs rather than the larger-model settings reported in prior work.

Table 6: Shared configurations for the matched ImageNet 256×256 experiments. Parameter count refers to the trainable DiT backbone and excludes the frozen autoencoder. The learning rate includes a 10k-update warm-up.

Configuration	Value
Network architecture and data	
Backbone	DiT-B/2
Backbone parameters (M)	approximately 131
Transformer depth / hidden dimension	12 / 768
Attention heads / latent patch size	12 / 2×2
MLP expansion ratio	4
Image resolution / latent resolution	256×256 / 32×32
Latent channels / conditioning classes	4 / 1,000
Frozen autoencoder	sd-vae-ft-ema
Training hyperparameters	
Training updates / fine-tuning	400k / none
Global batch size / training seed	256 / 1
Optimizer	Adam
Peak and post-warm-up learning rate	10^{-4}
Learning-rate schedule	10k linear warm-up, then constant
Adam (β_1, β_2) / ϵ	$(0.9, 0.95)$ / 10^{-8}
Weight decay / dropout	0 / 0
Gradient clipping norm	16
Training autocast precision	bfloat16
EMA half-life (updates) / ramp-up ratio	6,931 / 0.05
EMA update frequency	every update
Adaptive loss denominator offset	10^{-3}
Guidance $(\omega_{\mathrm{CFG}}, \kappa_{\mathrm{CFG}})$	$(1, 0)$ (disabled)
Class-label dropout probability	0
Generation and evaluation	
Generation checkpoint	400k-update EMA
Generated samples per metric evaluation	50,000
Noise standard deviation / time shift	1 / 1
Inference guidance scale	1
One-NFE time grid	$1 \to 0$
Two-NFE time grid	$1 \to 0.5 \to 0$
Two-NFE samplers	ODE and consistency
Feature metrics	FID; DINO-FID (DINOv2-ViT-L/14)

For EbbFlow and ablations (b)–(e) in Table 2, two GPUs process 128 samples each; the matched MeanFlow run uses four GPUs with 64 samples each. These runs retain the same global batch of 256. All ablations use the shared settings in Table 6 and $\alpha = 0$. Their changes to allocation, interval law, and horizon are listed in Table 5; active allocation and horizon schedules use the EbbFlow parameters in Table 7. The static allocation 0.82525 approximates the time-averaged scheduled allocation. Each training configuration has one seed. The AlphaFlow main-table row averages two metric evaluations of the same checkpoint.

Table 7: Temporal-sampler configurations for the matched baselines and EbbFlow. LN denotes $\mathrm{LogitNormal}(-0.4, 1.0^2)$; MM denotes min–max ordering of two independent draws, and LA denotes Equation 10. Dashes indicate an inactive schedule.

Configuration	MeanFlow	AlphaFlow	EbbFlow
Diagonal timestep law	LN	LN	LN
Interior coordinate law	MM (LN)	MM (LN)	LA (LN)
Diagonal fraction ρ_k	0.75	0.50	$1 \to 0.65$
Allocation schedule	constant	constant	sigmoid
Allocation start / end (updates)	—	—	0 / 400k
Allocation steepness γ_ρ	—	—	20
Allocation endpoint clamp	—	—	0.005
Interval horizon $h_{\max}$	1	1	$0.2 \to 1$
Horizon start / end (updates)	—	—	100k / 350k
Horizon power	—	—	2
Target parameter α	0	$1 \to 0$	0
Target schedule	constant	sigmoid	constant
Target start / end (updates)	—	0 / 400k	—
Target steepness / endpoint clamp	—	25 / 0.005	—

E QUALITATIVE SAMPLES

Figures 5 and 6 show samples from the 400k-update EMA checkpoint of the EbbFlow configuration in Table 7, which attains the lowest two-NFE ODE FID among the reported matched configurations (36.02). For each class, we use all seeds 0–15, without selecting individual outputs, and keep the class and initial noise fixed between one- and two-NFE generation. Within each sampler block, seeds appear in row-major order. Both samplers use guidance scale one; two-NFE ODE uses intermediate time 0.5. The full-horizon left-anchor control has the better one-NFE FID; these figures illustrate EbbFlow's joint curriculum rather than that separate control.

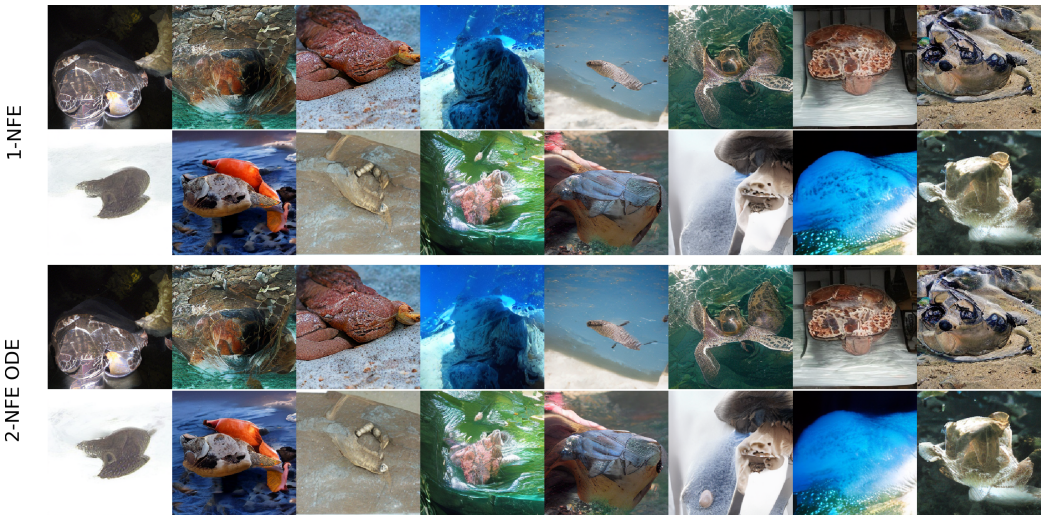


Figure 5: Uncurated EbbFlow samples: ImageNet class 33 (loggerhead), seeds 0–15, 400k EMA. Top: one NFE; bottom: two-NFE ODE.



Figure 6: Uncurated EbbFlow samples: ImageNet class 88 (macaw), seeds 0–15, 400k EMA. Top: one NFE; bottom: two-NFE ODE.