SAME FIXED POINT, DIFFERENT LEARNING DYNAMICS: ERROR TRANSPORT IN FINITE-STEP FLOW LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Training objectives for few-step generation can share a solution yet differ in the error feedback they supply during learning. For a frozen teacher, integrated endpoint labels stay fixed, whereas stopped differential labels return a directional derivative of the student's error. We show that an order-k horizon error receives diagonal residual gain $k + 1$, with material transport coupling different orders. These unequal gains let one global step size contract some modes and destabilize others in idealized target-space learning. In image networks, the kernel and AdamW reshape these gradient gains before they reach the learned function. Paired training across three image domains favors endpoint supervision; at five sampling steps, CIFAR-10 FID is 7.590 versus 21.133 for stopped-JVP training after 3,000 updates. Endpoint's advantage persists at equal measured time, including label production on shared, reusable training states. Choosing a target thus also chooses how different horizon orders are corrected.

1 INTRODUCTION

Few-step generation asks a network to predict a finite transition in one call. Training this network requires a choice of supervision: integrate a teacher trajectory to obtain a fixed endpoint, or use a differential identity to construct a label from the current prediction. Both can describe the same transition at convergence. Their agreement at the solution leaves open a practical question: how does each target teach a network that is still wrong?[1]

This choice affects both optimization and computation. Integrated labels cost teacher evaluations but can be reused. Differential labels avoid the full integration and require repeated student differentiation. Comparing these costs alone misses the feedback supplied by the labels. Recomputing a target returns part of the student's error to its next update, potentially changing which parts of the field are learned quickly or overshot.

Let $\Phi^{\mathrm{T}}_{t\to t+h}(x)$ be the flow map of a frozen teacher. We define

$$u^{\mathrm{T}}_h(t,x) = \frac{\Phi^{\mathrm{T}}_{t\to t+h}(x) - x}{h}. \tag{1}$$

The average-velocity field turns one Euler update into the teacher map: $x + hu^{\mathrm{T}}_h(t,x) = \Phi^{\mathrm{T}}_{t\to t+h}(x)$. Multiplication by h converts its prediction error into the error of the executed transition.

Integrated-velocity distillation supplies fixed displacement labels (Hu et al., 2026; Wang et al., 2026). MeanFlow represents the same average velocity through a differential identity, evaluated with the current network's Jacobian–vector product (JVP) (Geng et al., 2025). We compare these targets using one frozen teacher for both the endpoints and the derivative direction. Architecture, teacher states, and desired field remain shared. Their difference is whether the label returns a directional derivative of the student's current error.

Prior analyses connect network-dependent targets to gradient bias and unstable fixed-point learning (Hu et al., 2026; Kim et al., 2026b). We resolve this difference by horizon order. The stopped label

[1] Code and reproduction instructions:
https://anonymous.4open.science/r/error-transport-review-5FF9/.

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.

returns an order-k error with diagonal coefficient $-k$, while material derivatives transfer error into higher orders. The resulting residual gain $k+1$ changes the relative learning steps of different orders. It can accelerate one mode while making another oscillate. This order-dependent conditioning is the quantity we analyze and measure before and after neural optimization.

The network Jacobian and optimizer transform this target-side feedback into a function update. In image networks, they weaken the integer gains while preserving their ordering in injected-mode measurements. Paired image training links the choice of supervision to generation quality and learning cost. Agreement at the solution thus coexists with substantially different routes toward it.

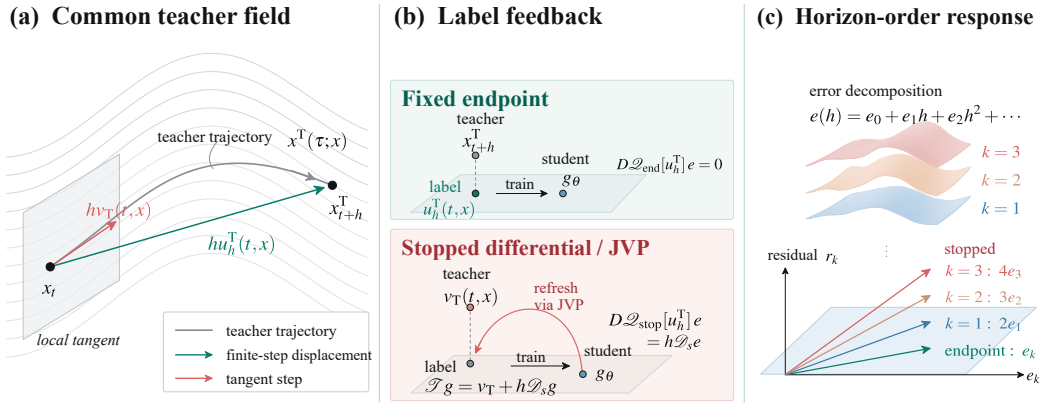

Figure 1: **Same teacher field, different supervision. a,** The average velocity reaches the teacher endpoint in one step. **b,** Endpoint labels stay fixed; stopped differential labels return a derivative of the student's error at each refresh. **c,** The formal horizon modes have diagonal residual gains $k+1$, compared with one for endpoint regression. The sheets illustrate modes; material transport couples their coefficients, and neural optimization acts after this target-side response.

2 Related Work

Learning finite-time flow maps. Probability-flow ODEs, Flow Matching, and Rectified Flow learn instantaneous velocities that numerical solvers integrate (Song et al., 2021; Lipman et al., 2023; Liu et al., 2023). Distillation and consistency models learn longer transitions directly (Salimans & Ho, 2022; Song et al., 2023). Consistency Trajectory Models and flow-map methods extend these transitions to pairs of times (Kim et al., 2024; Boffi et al., 2025a; Sabour et al., 2025). MeanFlow parameterizes their average velocity through a differential identity (Geng et al., 2025). CMT and Int-MeanFlow construct fixed velocity labels by integrating teacher trajectories (Hu et al., 2026; Wang et al., 2026). These constructions supply the integrated and differential targets studied here. Mean Flow Distillation uses integrated teacher and auxiliary fields to guide a generator's distribution (Zhao et al., 2026).

Fixed points and training updates. Equivalent population or PDE solutions connect several consistency formulations (Lai et al., 2023). Eulerian and Lagrangian self-distillation differ in gradient stability (Boffi et al., 2025b); continuous-time consistency benefits from stabilization (Lu & Song, 2025). CMT relates network-dependent pseudo-targets to discrepancies from the oracle regression gradient (Hu et al., 2026). Stabilizing Consistency Training analyzes how stopped objectives can retain the desired flow-map fixed point while changing convergence behavior (Kim et al., 2026b). Our horizon-order analysis resolves the returned error into diagonal gains and material transport, then measures their effect through neural optimization.

Optimizing average-velocity models. Improved MeanFlow rewrites the loss in instantaneous-velocity space and uses a learned marginal velocity as the JVP direction (Geng et al., 2026). AlphaFlow separates trajectory matching from consistency, measures their gradient conflict, and schedules a finite consistency-step ratio (Zhang et al., 2026). Other analyses study label variance and the order in which short- and long-interval velocities are learned (Lu & Wang, 2026; Kim et al.,

2026a). Holding the teacher direction fixed lets us distinguish the response of existing horizon errors from changes in the desired field.

3 Learning the Same Finite-Step Field

Let a frozen teacher solve $\dot{X}_t = v_{\mathrm{T}}(t, X_t)$ on $\mathbb{R}^d$, with flow map $\Phi^{\mathrm{T}}_{t\to s}(x)$. On a deployment grid $\mathcal{G} = (t_0, \dots, t_N)$, write $h_i = t_{i+1} - t_i \neq 0$. Signed steps cover increasing flow time and decreasing noise coordinates. Euler sampling applies $\mathcal{E}^{f_\theta}_{h_i}(t_i, x) = x + h_i f_\theta(t_i, x, h_i)$.

The average velocity in Eq. (1) is the field this sampler must learn. Fix the right endpoint $s = t + h$ and move the left endpoint along a teacher characteristic. The corresponding derivative and label are

$$\mathcal{D}_s = v_{\mathrm{T}} \cdot \nabla_x + \partial_t - \partial_h, \qquad \mathcal{T} g = v_{\mathrm{T}} + h \mathcal{D}_s g. \tag{2}$$

Differentiating the teacher endpoint gives $\mathcal{T} u^{\mathrm{T}}_h = u^{\mathrm{T}}_h$. We compare the two label maps

$$\mathcal{Q}_{\mathrm{end}}(g) = u^{\mathrm{T}}_h, \qquad \mathcal{Q}_{\mathrm{stop}}(g) = \mathcal{T} g. \tag{3}$$

Both fix the same average-velocity field. Endpoint regression uses the integrated teacher label; stopped-JVP regression evaluates $\mathcal{D}_s f_\theta$ with direction $(1, v_{\mathrm{T}}, -1)$ in (t, x, h). Both losses regress f_θ against a detached label: $\mathcal{L}_A(\theta) = \mathbb{E}_\nu \| f_\theta - \mathrm{sg}[\mathcal{Q}_A(f_\theta)] \|_2^2$, where $A \in \{\mathrm{end}, \mathrm{stop}\}$ and ν samples (t, x, h). Stop-gradient holds the label constant during an optimizer step. The stopped label is recomputed at the next step, using the updated network. This across-step dependence is the difference studied here.

We distinguish the mathematical flow Φ^{T} from the numerical map $\widehat{\Phi}^{\mathrm{T}}$ used to build labels. A high-accuracy teacher rollout supplies

$$X^{\mathrm{T}}_{i+1} = \widehat{\Phi}^{\mathrm{T}}_{t_i \to t_{i+1}}(X^{\mathrm{T}}_i), \tag{4}$$

$$\widehat{u}^{\mathrm{T}}_{h_i}(t_i, X^{\mathrm{T}}_i) = \frac{X^{\mathrm{T}}_{i+1} - X^{\mathrm{T}}_i}{h_i}. \tag{5}$$

The numerical endpoint replaces u^{T}_h in the endpoint loss. Paired comparisons share the teacher and training states; Section 5 specifies the matched training budgets and independent rate selection. Appendix A.1 specifies the construction and evaluation of both labels.

The instantaneous velocity v_{T} is the $h = 0$ boundary of this field and initializes both image students. Its finite-step correction provides the local prediction in Section 4.2. A separate tangent-supervision comparison measures the effect of learning that correction.

4 Target-Error Transport in Horizon Space

We assume a unique smooth teacher flow and finite squared risks. Appendix D gives the proofs and the propagation of field error through a composed sampler. Here we study how the label transforms that error before an optimizer step.

4.1 A stopped target transports the student's error

Subtracting $\mathcal{T} u^{\mathrm{T}}_h = u^{\mathrm{T}}_h$ from the recomputed label gives, for every smooth predictor g,

$$\mathcal{T} g - u^{\mathrm{T}}_h = h \mathcal{D}_s (g - u^{\mathrm{T}}_h). \tag{6}$$

With a frozen teacher direction, the iMF interval residual $g - h\,\mathrm{sg}[\mathcal{D}_s g] - v_{\mathrm{T}}$ produces the same update. The full iMF objective also learns the velocity used inside this derivative (Geng et al., 2026).

The two label maps in Eq. (3) have different first variations around their common field:

$$D\mathcal{Q}_{\mathrm{end}}[u^{\mathrm{T}}_h] e = 0, \qquad D\mathcal{Q}_{\mathrm{stop}}[u^{\mathrm{T}}_h] e = h \mathcal{D}_s e.$$

Here $D\mathcal{Q}$ measures sensitivity across label refreshes. It describes how an error in the current prediction enters the next label.

We hold the external endpoint field fixed and scale only the residual feedback:

$$\mathcal{T}_\lambda g = u_h^{\mathrm{T}} + \lambda h \mathcal{D}_s (g - u_h^{\mathrm{T}}) = (1-\lambda) u_h^{\mathrm{T}} + \lambda \mathcal{T} g, \qquad 0 \le \lambda \le 1. \tag{7}$$

On characteristic intervals reaching $h = 0$, boundedness selects this fixed point uniquely. For $\lambda > 0$, a fixed-point error obeys $e = \lambda h \mathcal{D}_s e$, giving $e = C|h|^{-1/\lambda}$ along each fixed-endpoint characteristic. Bounded error forces $C = 0$.

The idealized relaxation $g^+ = (1-\alpha)g + \alpha \mathcal{T}_\lambda g$ updates its error by $e^+ = [I - \alpha(I - \lambda h \mathcal{D}_s)]e$. At $\lambda = 0$, every direction receives the same relaxation. Positive λ returns transported error before the model kernel and optimizer act. We resolve this target-side response by horizon degree.

4.2 Tangent boundary and factor-two initialization

Define $a_{\mathrm{T}} = \partial_t v_{\mathrm{T}} + D_x v_{\mathrm{T}} v_{\mathrm{T}}$. At tangent initialization $g_0 = v_{\mathrm{T}}$, the first labels obey

$$u_h^{\mathrm{T}}(t,x) = v_{\mathrm{T}}(t,x) + \tfrac{h}{2} a_{\mathrm{T}}(t,x) + O(h^2), \tag{8}$$

$$\mathcal{T} g_0(t,x,h) = v_{\mathrm{T}}(t,x) + h a_{\mathrm{T}}(t,x). \tag{9}$$

The stopped label doubles the leading correction $\Delta_h^{\mathrm{T}} = u_h^{\mathrm{T}} - v_{\mathrm{T}}$. The initialization and its first refreshed label lie on opposite sides of the endpoint field, with leading errors $-h a_{\mathrm{T}}/2$ and $+h a_{\mathrm{T}}/2$. Full endpoint relaxation removes that error; full stopped relaxation reflects it. This is the $k = 1$ response in label space and supplies a signed, parameter-free prediction for the local experiments.

4.3 Formal horizon-jet response

Let V be a finite-dimensional coefficient space and write a degree-K formal jet as $e(h) = \sum_{k=0}^{K} h^k e_k + O(h^{K+1})$, with $e_k \in V$. The projection Π_K retains the coefficient tuple in $\mathcal{J}_K(V) = V^{K+1}$.

Theorem 1 (Formal horizon-jet response)**.** *Suppose $\mathcal{L}e = h\{\mathcal{A}(h)e - \partial_h e\}$, where $\mathcal{A}(h) = \sum_{j \ge 0} h^j A_j$ and each $A_j : V \to V$ is linear. Modulo h^{K+1}, the induced map on $\mathcal{J}_K(V)$ is block triangular with diagonal blocks $0, -I, -2I, \dots, -KI$. Its formal horizon-jet spectrum is*

$$\operatorname{spec}(\Pi_K \mathcal{L}) = \{0, -1, -2, \dots, -K\}. \tag{10}$$

If $e(0) = 0$, the constant block is absent. The idealized relaxed update $g^+ = (1-\alpha)g + \alpha \mathcal{T}_\lambda g$ multiplies the diagonal order-k response by $1 - \alpha(1 + \lambda k)$. The finite triangular update is asymptotically stable through order K when

$$0 < \alpha < \frac{2}{1 + \lambda K}. \tag{11}$$

For the frozen-teacher MeanFlow operator, $\mathcal{A} = v_{\mathrm{T}} \cdot \nabla_x + \partial_t$. Projecting its coefficient operators onto a finite basis gives the maps A_j above. The degree-raising part $h\mathcal{A}(h)e$ can mix lower orders into higher ones and create transient growth without changing the diagonal eigenvalues.

For a pure mode $h^k e_k$, horizon differentiation returns $-k h^k e_k$, whereas $h\mathcal{A}(h)e$ begins at degree $k+1$. The integers in Eq. (10) count horizon degree, not spatial frequency or model depth.

The two learning problems therefore induce different coefficient updates:

$$\begin{aligned} e^+_{k,\mathrm{end}} &= (1-\alpha) e_k, \\ e^+_{k,\mathrm{stop}} &= \{1 - \alpha(1 + \lambda k)\} e_k + \alpha\lambda \sum_{j+\ell=k-1} A_j e_\ell. \end{aligned} \tag{12}$$

The sum vanishes for a pure order-k perturbation. At $\lambda = 1$, its diagonal update moves the linear mode twice as far, the quadratic mode three times as far, and so on. A larger gain can improve contraction: at $\alpha = 0.1$, the linear-mode multiplier is 0.8 under stopped supervision and 0.9 under endpoint supervision. The difficulty is their relative scale across orders. A step that contracts low-order error can overshoot higher-order structure. Reducing the global step reduces every mode's update, including low-order modes that could tolerate a larger step. If the model enforces $e(0) = 0$, the constant mode is absent; otherwise it also supplies higher-order error through the sum.

From label feedback to a learned function. The label response determines a residual, not a parameter step. The network pulls that residual back through its Jacobian, and the optimizer rescales and combines the resulting parameter directions. At an attainable fixed point $f_{\theta^\star} = g^\star$, a stopped squared-loss update with a fixed preconditioner B has differential

$$\delta f^+ = [I - \eta K_{\theta,B}\{I - D\mathcal{Q}[g^\star]\}]\,\delta f, \qquad K_{\theta,B} = J_\theta B J_\theta^*, \tag{13}$$

This tangent-kernel description (Jacot et al., 2018) uses J_θ to map parameter perturbations to functions. Its adjoint uses the training measure, and η absorbs the loss factor. For error e and residual $r = e - h\mathcal{D}_s e$, first-order training-risk change has sign $-\langle J_\theta^* e, BJ_\theta^* r\rangle$. The kernel $K_{\theta,B}$ can mix horizon modes: the target's integer gains need not remain integer ratios in the learned function. Unequal gains on nonorthogonal modal gradients can rotate updates away from descent, even with stable linearized dynamics. Our probes measure AdamW's full step, including its evolving moments (Appendix D.3).

The same operator can transport unrepresentable teacher error into a trainable direction, shifting a finite model's fitted solution. Appendix B.1 derives this projected equilibrium shift and tests it with a quartic teacher and cubic student. Our attainable neural teacher removes this approximation effect and isolates the learning dynamics.

4.4 Finite shifts and the differential limit

The same calculation applies to the fixed-direction finite consistency target of AlphaFlow (Zhang et al., 2026). With step ratio a, write $\mathcal{Q}_a g = av_{\mathrm{T}} + (1-a)g(t+ah, x+ahv_{\mathrm{T}}, (1-a)h)$. Its loss weight $1/a$ gives the residual $\mathcal{R}_a(g) = (g - \mathcal{Q}_a g)/a$, whose order-$k$ diagonal gain is

$$\beta_k(a) = \frac{1-(1-a)^{k+1}}{a} = \sum_{j=0}^{k}(1-a)^j, \qquad \beta_k(1) = 1, \quad \lim_{a\to 0}\beta_k(a) = k+1. \tag{14}$$

Larger Euler shifts soften the high-order response and introduce forcing: $b_a = \mathcal{Q}_a u_h^{\mathrm{T}} - u_h^{\mathrm{T}} = -a^2 h\mathcal{A}v_{\mathrm{T}}/2 + O(h^2)$, with $\mathcal{A} = \partial_t + v_{\mathrm{T}}\cdot\nabla_x$. At $a = 1$ the label is the instantaneous velocity; as $a \to 0$ the normalized residual recovers the differential identity. This Euler shift changes response and target bias together, whereas λ in Eq. (7) varies feedback around a fixed field. Appendix B.3 separates the response from the teacher-integration error that produces forcing.

5 Experimental Setup

Polynomial and 32-parameter tanh controls use five paired initialization seeds and attainable teachers. The neural model trains for 30,000 Adam updates, with an independent state–horizon grid for evaluation.

Image experiments use released EMA teachers on CIFAR-10 (32×32), FFHQ-64, and AFHQv2-64. The endpoint–stopped-JVP comparison uses three independent training seeds per domain and 3,000 updates, with a checkpoint at 1,000. Each pair shares its architecture, initialization, teacher states, minibatch sequence, optimizer family, and generation noise. Target-specific learning rates are chosen on separate pilot seeds. CIFAR-10 uses continuous start times and horizons; the 64×64 models use eight teacher intervals. The primary sampling budgets are NFE 5 and NFE 8, respectively.

Each cross-domain seed uses two common generation streams: 50,000 images per stream for CIFAR-10 CleanFID and 10,000 for 64×64 EDM FID. We average streams within each seed and compute sample SD and paired 95% intervals across seeds. Appendix A.1 specifies probe sizes, learning-rate selection, and costs.

The cost comparison uses a shared 1,000-row CIFAR-10 state–anchor dataset with $h \in [1/128, 1/2]$ and equal-weight endpoint supervision at $h_0 = 1/128$. Three initialization/order seeds compare label precision at 1,000 updates; three further seeds compare targets at 300 and 600 seconds after independent pilot selection of each target's learning rate. The clock includes on-demand labels, transfers, and student updates. Each model uses one shared 50,000-image stream at NFE 5 and 7 (Appendix A.3).

6 RESULTS

6.1 HORIZON ORDER GOVERNS TARGET-SPACE STABILITY

Higher-order errors reach instability at smaller learning steps, even when the teacher is exactly representable (Figure 2). Injected horizon modes have measured slopes $-1, -2, -3$ and cross unit response at $\alpha = 2/(k+1)$. For a cubic student, feedback weights $\lambda = 0, 0.5, 1$ have empirical stability brackets containing the predicted ceilings $2, 0.8, 0.5$. At $\alpha = 0.49$, the endpoint cubic mode contracts with multiplier 0.51, while the stopped mode alternates with multiplier -0.96. Increasing the step to 0.51 makes only the stopped arm unstable.

The difference persists in nonlinear learning. After 1,000 Adam updates from perturbed teacher weights, endpoint and stopped training have mean relative held-out RMS errors 0.00440 and 0.01836 across five seeds. At 30,000 updates, every arm is within 5.19×10^{-5} of the same teacher. The function kernel changes by 7.3–10.6% over training. Both targets approach the same representable function along different learning trajectories.

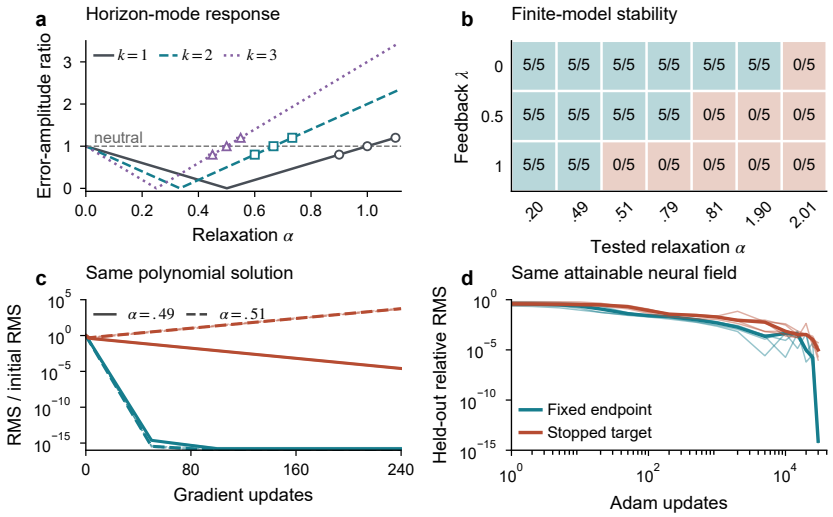

Figure 2: **Different learning trajectories toward an attainable common solution. a,** Predicted $|1-\alpha(k+1)|$ (lines) and injected-mode measurements (symbols, 95% intervals). **b,** Stable polynomial runs among five seeds at each (λ, α). **c,** Polynomial gradient descent at $\alpha = 0.49$ (solid) and 0.51 (dashed). **d,** Adam training of a tanh model with an exactly representable teacher. Thin traces in c–d are five seeds; thick traces are means in c and medians in d.

Corrections at small and finite horizons. The first refreshed label nearly doubles the endpoint correction: projected JVP factors are 2.006, 1.947, and 1.931 on CIFAR-10, FFHQ-64, and AFHQv2-64 (Figure 3). Direction cosines of 0.986–0.999 and horizon exponents of 0.968–1.065 recover the predicted direction and scaling. Refining endpoint integration from two to sixteen RK4 microsteps preserves the 64×64 projection factors at 1.93–1.95.

Moving away from the differential limit weakens high-order feedback and introduces teacher-integration error. In an exact nonnormal affine flow with state-dependent perturbations (Figure 4), recomputed targets recover the predicted $\beta_k(a)$ curves. Measured forcing scales as $h^{0.997–1.000}$ and $a^{2.000}$, following the predicted $a^2 h$ law. The shift ratio thus changes both the response to student error and the forcing present at the teacher field.

6.2 NEURAL OPTIMIZATION RESHAPES THE TARGET RESPONSE

Neural optimization weakens the target's order-dependent amplification (Figure 5b). At two trained checkpoints, injected orders $k = 1, 2, 3$ give median parameter-gradient ratios $2.00, 3.00, 4.02$, close to $1 + k$. After the network kernel and AdamW act, the held-out function ratios are $1.51, 2.05, 2.53$.

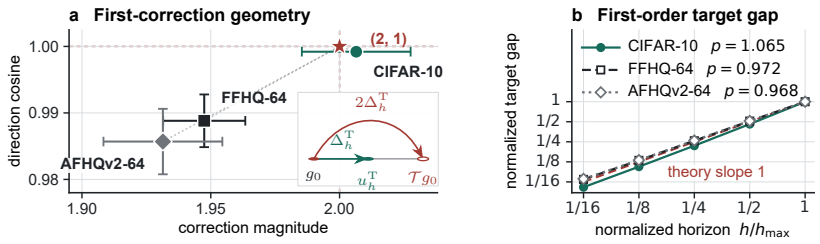

Figure 3: **First-correction geometry recovers the factor-two law. a,** Correction magnitudes and direction cosines approach $(2, 1)$ (95% intervals); the inset shows reflection around u_h^{T}. **b,** Target gaps follow the slope-one guide; the legend gives fitted horizon exponents.

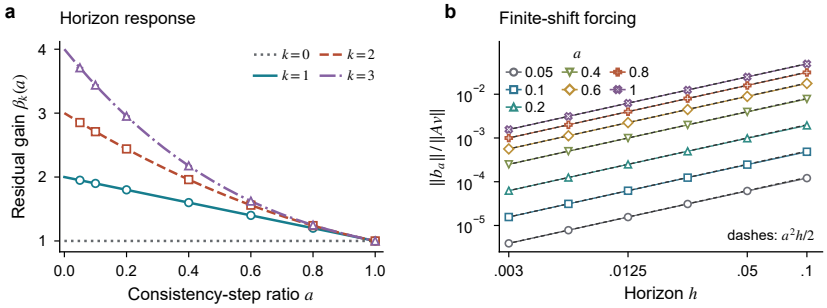

Figure 4: **Euler consistency shifts: response and integration error. a,** Predicted $\beta_k(a)$ (lines) and recomputed-target gains (markers, eight states per ratio and order); statewise ranges are smaller than markers and $k = 0$ is analytical. **b,** Acceleration-normalized forcing: eight states (faint), their mean (markers), and $a^2h/2$ (dashed). Both panels measure labels before neural optimization.

The probes add spatially constant output perturbations, isolating the horizon derivative in the actual image network.

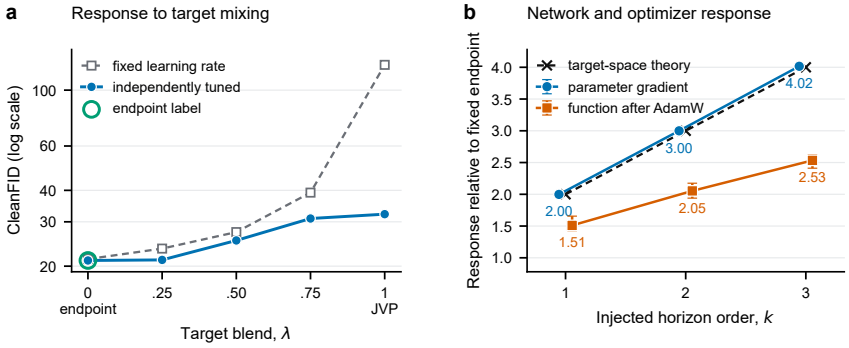

Figure 5: **Image learning and update responses under target feedback. a,** NFE-5 CleanFID after 1,000 matched updates: independently tuned rates (blue), a fixed-rate control (gray), and endpoint supervision (green). **b,** Injected-mode response in parameter gradients (blue) and held-out functions after AdamW (orange). Medians and interquartile ranges pool eight perturbations across the 1,000- and 3,000-update checkpoints.

Natural errors also require material transport. Along fixed-endpoint teacher paths, $\mathrm{d}e/\mathrm{d}h = -\mathcal{D}_s e$ includes changing t and x, unlike the fixed-state jet. At twelve checkpoints, applying this operator to cubic training-batch error fits reduces gradient reconstruction error by 52–59% relative to positive scalar rescaling. Gradient cosines are 0.900–0.925 across three seeds (Appendix Table 18). Fixed-state measurements separate horizon differentiation from state transport (Appendix C.4; Table 15).

What target feedback changes during learning. The returned error can oppose endpoint descent. Across sixteen batches at three 3,000-update JVP checkpoints, horizon feedback cancels most or all of the positive endpoint gradient work. State transport further opposes descent (Appendix Table 20). After AdamW, JVP gives higher cross-batch field error than endpoint in two models and a near-zero difference in the third.

Close gradient alignment can nevertheless mispredict descent. Cubic reconstruction gives projection-sign balanced accuracies of 85.7%, 75.0%, and 56.7% across seeds: the remaining error can change the component along field-error descent (Appendix Tables 18–20).

Scaling the full feedback with λ preserves the teacher field while changing both transport terms. With separately tuned rates, 1,000-update NFE-5 CleanFID rises from 21.068 to 32.172 as λ increases from zero to one (Figure 5a; Appendix C.5).

Field error and differential residual can rank models differently. At 3,000 FFHQ-64 updates, three paired seeds give mean field RMSE 0.01155 for endpoint and 0.03401 for JVP, but differential residuals 0.03016 and 0.02131. JVP fits its differential constraint more closely while remaining farther from the integrated field: weighting and mixing error components changes their norm's ranking. CIFAR-10 favors endpoint on both norms, while AFHQv2-64 reproduces the FFHQ ordering (Appendix Table 13). Continuous FFHQ intervals preserve this distinction (Appendix C.1).

6.3 Generation quality and learning cost

Fixed endpoint supervision gives lower FID across three independently trained pairs per domain (Table 1); paired 95% intervals exclude zero (Appendix Table 12). At NFE 7, the same CIFAR-10 checkpoints give 7.741 versus 26.795 (Appendix Table 12). Figure 7 shows an independent pair's same-noise outputs at both budgets.

Table 1: **Fixed endpoint and stopped-JVP supervision across image domains.** FID after 3,000 updates, mean ± sample SD over three training seeds. Each seed averages two common generation streams: 50,000 images each for CIFAR-10, and 10,000 each for FFHQ-64 and AFHQv2-64.

		FID ↓	
Dataset	NFE	Fixed endpoint	Stopped JVP
CIFAR-10	5	$\mathbf{7.590 \pm 0.098}$	21.133 ± 2.979
FFHQ-64	8	$\mathbf{5.550 \pm 0.033}$	33.726 ± 0.251
AFHQv2-64	8	$\mathbf{3.948 \pm 0.058}$	14.921 ± 0.337

Recomputing teacher endpoints on either student's rollout states preserves endpoint training's advantage in local and composed MSE and LPIPS (Zhang et al., 2018) (Figure 6; Appendix Table 24).

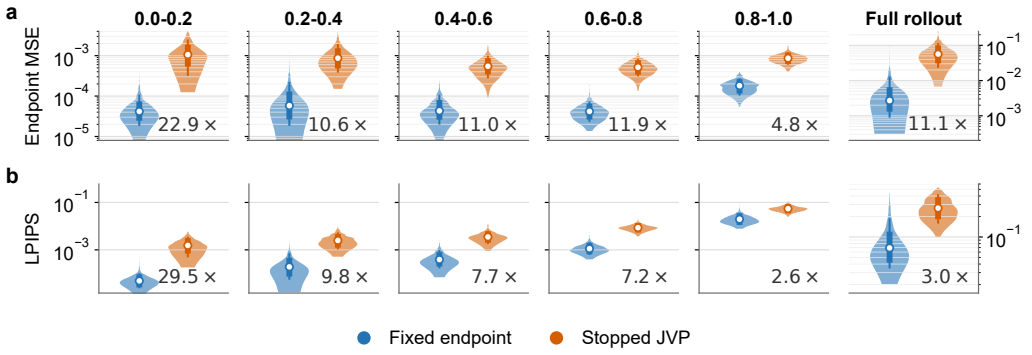

Figure 6: **The target gap persists along the full student rollout.** Per-noise distributions over 512 fixed inputs compare states visited by either student. Endpoint MSE (**a**) and LPIPS (**b**) cover five intervals and the final composition; annotations give stopped-JVP to fixed-endpoint mean ratios.

Two-step RK4 uses eight teacher calls per label, raising mean FID by 0.153/0.058 at NFE 5/7 after 1,000 updates (Table 2). The equal-time comparison times interior labels and training on 1,000

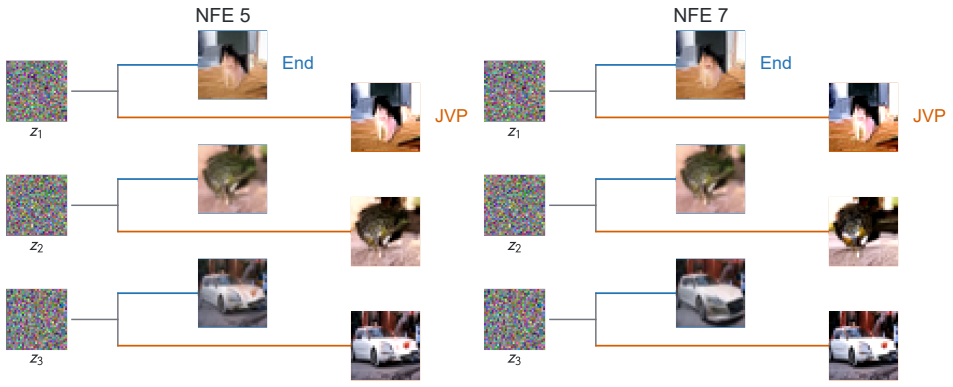

Figure 7: **Same-noise samples from fixed endpoint and stopped-JVP training.** An independent 3,000-update CIFAR-10 pair uses the first three Gaussian inputs in stream order at NFE 5 and 7. Each budget starts independently from the same noise; colored branches distinguish the two learned fields.

reusable batch rows, starting from shared states and boundary labels. Both arms cache teacher quantities; the JVP target is recomputed at each update. Independent rate selection favors endpoint at both budgets, despite fewer updates at 300 seconds. The 600-second NFE-5/7 gaps are 1.063/1.292 FID. Common-rate and half-rate controls appear in Appendix Tables 8 and 9.

Table 2: **Supervision precision and learning cost. Top:** Endpoint after 1,000 updates; teacher evaluations count images. **Bottom:** Independently selected rates (both 2×10^{-5}); A100 time includes interior labels, transfers and updates on shared states and boundary labels. Endpoint uses two-step RK4. Each panel uses three paired seeds and 50,000 CIFAR-10 images per FID; uncertainty is sample SD. Updates and label times are means (Appendix A.3).

Supervision precision: 1,000 updates

Endpoint label	Interior-label teacher evaluations (M)	FID ↓ NFE 5	FID ↓ NFE 7
Reference integration	8.010624	$\mathbf{8.856 \pm 0.230}$	$\mathbf{7.498 \pm 0.305}$
Two-step RK4	0.256000	9.009 ± 0.201	7.556 ± 0.268

Target choice at equal measured time

Time (s)	Target	Updates	Label time (s)	Field RMSE ↓	FID ↓ NFE 5	FID ↓ NFE 7
300	Endpoint	636.7	126.9	0.0575 ± 0.0025	$\mathbf{8.314 \pm 0.305}$	$\mathbf{7.085 \pm 0.173}$
	Stopped JVP	686.0	17.8	0.0674 ± 0.0026	12.254 ± 0.721	11.726 ± 0.816
600	Endpoint	1478.7	199.2	0.0480 ± 0.0006	$\mathbf{7.093 \pm 0.058}$	$\mathbf{6.226 \pm 0.095}$
	Stopped JVP	1396.0	25.9	0.0603 ± 0.0002	8.155 ± 0.460	7.518 ± 0.401

Supervision within a complete generation method. Coupling supervision to a denoising map and re-noising sampler changes the generation comparison: native two-call consistency attains FID 5.960 versus deterministic endpoint composition's 6.716 at equal student-forward counts (Appendix C.7). This advantage belongs to the combined choice of predicted object and sampler.

7 Conclusion

Supervision determines both a desired field and the feedback used to learn it. Recomputed differential labels return horizon errors with unequal gains; the network and optimizer transform those gains into function updates. The resulting order-dependent conditioning explains how learning trajectories can diverge around a common solution. In paired image training with reusable teacher states, fixed endpoint labels give better generation quality at both matched updates and measured time. A target's fixed point specifies what is learned; its error transport helps determine how learning proceeds.

AI USE STATEMENT

Generative AI tools assisted with manuscript editing and readability, figure refinement, and aspects of research planning, coding, and analysis. The authors take responsibility for the final manuscript and reported results.

REFERENCES

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*, 2025a. URL https://openreview.net/forum?id=cqDH0e6ak2.

Nicholas M. Boffi, Michael S. Albergo, and Eric Vanden-Eijnden. How to build a consistency model: Learning flow maps via self-distillation. In *Advances in Neural Information Processing Systems*, volume 38, 2025b. doi: 10.52202/085713-1121. URL https://arxiv.org/abs/2505.18825.

Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. StarGAN v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. URL https://arxiv.org/abs/1912.01865.

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, 2025. doi: 10.52202/085713-2534. URL https://arxiv.org/abs/2505.13447.

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/papers/Geng_Improved_Mean_Flows_On_the_Challenges_of_Fastforward_Generative_Models_CVPR_2026_paper.pdf.

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://arxiv.org/abs/1706.08500.

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://openreview.net/forum?id=2B8GkGTgmY.

Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*, volume 31, 2018. URL https://arxiv.org/abs/1806.07572.

Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4401–4410, 2019. URL https://arxiv.org/abs/1812.04948.

Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. In *Advances in Neural Information Processing Systems*, volume 34, 2021. URL https://arxiv.org/abs/2106.12423.

Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Advances in Neural Information Processing Systems*, volume 35, 2022. URL https://arxiv.org/abs/2206.00364.

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://arxiv.org/abs/2310.02279.

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, 2026a. URL https://openaccess.thecvf.com/content/CVPR2026/html/Kim_Understanding_Accelerating_and_Improving_MeanFlow_Training_CVPR_2026_paper.html.

Youngjoong Kim, Duhoe Kim, Woosung Kim, and Jaesik Park. Stabilizing consistency training: A flow map analysis and self-distillation, 2026b. URL https://arxiv.org/abs/2601.22679.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015. URL https://arxiv.org/abs/1412.6980.

Alex Krizhevsky. 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.

Chieh-Hsin Lai, Yuhta Takida, Toshimitsu Uesaka, Naoki Murata, Yuki Mitsufuji, and Stefano Ermon. On the equivalence of consistency-type models: Consistency models, consistent diffusion models, and Fokker–Planck regularization. *arXiv preprint arXiv:2306.00367*, 2023. doi: 10.48550/arXiv.2306.00367. URL https://arxiv.org/abs/2306.00367.

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://arxiv.org/abs/2210.02747.

Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond. In *International Conference on Learning Representations*, 2020. URL https://arxiv.org/abs/1908.03265.

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://arxiv.org/abs/2209.03003.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL https://arxiv.org/abs/1711.05101.

Cheng Lu and Yang Song. Simplifying, stabilizing and scaling continuous-time consistency models. In *International Conference on Learning Representations*, 2025. URL https://arxiv.org/abs/2410.11081.

Juanwu Lu and Ziran Wang. On variance reduction in learning mean flows. *arXiv preprint arXiv:2605.09235*, 2026. doi: 10.48550/arXiv.2605.09235. URL https://arxiv.org/abs/2605.09235.

Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in GAN evaluation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. URL https://arxiv.org/abs/2104.11222.

Amirmojtaba Sabour, Sanja Fidler, and Karsten Kreis. Align your flow: Scaling continuous-time flow map distillation. In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-4900. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/d79ac139911df25d27f14bbf008deaee-Abstract-Conference.html.

Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*, 2022. URL https://arxiv.org/abs/2202.00512.

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://arxiv.org/abs/2011.13456.

Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 32211–32252, 2023. URL https://proceedings.mlr.press/v202/song23a.html.

Wei Wang, Rong Cao, Yi Guo, Zhengyang Chen, Kuan Chen, and Yuanyuan Huo. Int-MeanFlow: Few-step speech generation with integral velocity distillation. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 18317–18321, 2026. doi: 10.1109/ICASSP55912.2026.11462627. URL https://doi.org/10.1109/ICASSP55912.2026.11462627.

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.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. URL https://arxiv.org/abs/1801.03924.

An Zhao, Shengyuan Zhang, Zhongjian Sun, Yixiang Zhou, Zejian Li, Ling Yang, Tianrun Chen, and Lingyun Sun. Mean flow distillation: Robust and stable distillation for flow matching models. *arXiv preprint arXiv:2606.11155*, 2026. doi: 10.48550/arXiv.2606.11155. URL https://arxiv.org/abs/2606.11155.

A Additional Experimental Details

A.1 Paired endpoint and stopped-JVP training

The two objectives use the same frozen teacher and student parameterization. The datasets are CIFAR-10 (Krizhevsky, 2009), FFHQ (Karras et al., 2019), and AFHQ (Choi et al., 2020); we use the revised AFHQv2 release associated with Alias-Free GAN (Karras et al., 2021). The 64×64 teachers and evaluation implementation come from EDM (Karras et al., 2022). CIFAR-10 has 35,763,267 trainable parameters; each 64×64 model has 61,805,571. Zero-initialized horizon conditioning makes both students equal to their pretrained teacher before adaptation. We use AdamW (Loshchilov & Hutter, 2019) with $\beta = (0.9, 0.999)$, weight decay 0.01, and gradient-norm clipping at 1. Each loss is the mean squared error between the prediction and a detached label. Both models remain in evaluation mode during adaptation.

Table 3: **Paired image-training settings.** Learning rates are chosen before the three evaluation seeds. Every run uses 3,000 updates and a 100-update linear warmup.

	CIFAR-10	FFHQ-64	AFHQv2-64
Batch size	64	32	32
Endpoint peak learning rate	2×10^{-5}	2×10^{-5}	2×10^{-5}
Stopped-JVP peak learning rate	5×10^{-6}	2×10^{-5}	2×10^{-5}
Primary sampling NFE	5	8	8
Images per generation stream	50,000	10,000	10,000
Student arithmetic	BF16	FP32	FP32

After warmup, the learning rate decays by cosine to one tenth of its peak. CIFAR-10 reaches that rate at update 1,000 and holds it through update 3,000; the two 64×64 models decay through update 3,000. CIFAR-10 retains FP32 parameters, teacher labels, and loss accumulation under BF16 student autocasting. We evaluate raw checkpoint weights at updates 1,000 and 3,000. The final checkpoint and primary NFE are fixed before evaluation.

Teacher states and labels. Each domain has 12,288 teacher paths integrated with RK4 at NFE 64. The last 256 paths supply diagnostics; the remaining 12,032 supply training. The 64×64 models draw left states uniformly from their eight cached intervals. CIFAR-10 draws continuous starts and horizons according to Eqs. (16)–(19). Its start state, endpoint displacement, and teacher velocity are prepared once for each training sequence and reused by both objectives. The same sequence of states, times, and horizons enters every paired update.

Training seeds are 920101/920103/920107 for CIFAR-10, 921101/921103/921107 for FFHQ-64, and 922101/922103/922107 for AFHQv2-64. At each saved checkpoint, both objectives are evaluated on the same 128 held-out states. For d scalar output coordinates, define $\|r\|_{\mathcal{D}}^2 = (128d)^{-1} \sum_i \|r(z_i)\|^2$. We measure

$$E_{\text{field}} = \|g - \widehat{u}^{\mathrm{T}}\|_{\mathcal{D}}, \qquad E_{\text{diff}} = \|g - \mathcal{T}g\|_{\mathcal{D}}. \tag{15}$$

The first measures approximation of the teacher's integrated field; the second measures satisfaction of the differential identity at the sampled states.

Learning-rate selection. Stopped JVP uses four peak rates, 2×10^{-5} times $1, 1/2, 1/4, 1/8$, evaluated after 1,000 updates on independent pilot seeds. The respective 2,000-image FIDs are 126.001/206.318/32.172/57.194 on CIFAR-10, 42.136/44.394/48.191/54.975 on FFHQ-64, and 19.690/20.192/20.227/21.145 on AFHQv2-64. The selected multipliers are $1/4$, 1, and 1. Pilot training seeds are 190071, 280071, and 922011, respectively. Endpoint training retains its previously selected peak rate of 2×10^{-5}. Evaluation seeds are separate from these choices.

Generation and uncertainty. Generation seeds 924001 and 924007 are shared by the two objectives. FID (Heusel et al., 2017) on CIFAR-10 uses the CleanFID implementation (Parmar et al., 2022) in `legacy_tensorflow` mode against CIFAR-10 training-set statistics. FFHQ-64 and AFHQv2-64 use the NVLabs EDM Inception detector and dataset reference statistics. Each comparison holds the evaluator, reference, sample count, and sampling grid fixed. CIFAR-10 also has a secondary NFE-7 evaluation.

Let F_{ir}^{A} be FID for target A, training seed i, and generation stream r. We first form $\overline{F}_i^{A} = (F_{i1}^{A} + F_{i2}^{A})/2$. Arm summaries use the mean and sample SD of these three training-level values. For $\delta_i = \overline{F}_i^{\mathrm{JVP}} - \overline{F}_i^{\mathrm{end}}$, the paired interval is $\overline{\delta} \pm 4.30265\, s_\delta/\sqrt{3}$. The independent replication unit is a trained model pair.

Continuous labels and horizon parameterization. Continuous CIFAR-10 labels draw independent start fractions and spans:

$$R \sim \mathrm{Uniform}[1,4], \qquad S \sim \mathrm{Uniform}[0,1], \qquad T = S\left(1 - \frac{R}{8}\right). \tag{16}$$

The physical step is $H = R/8$, with $T + H \le 1$. From an independently sampled cached path, take X_A^{T} at $A = \lfloor 8T \rfloor/8$ and construct

$$X_T^{\mathrm{T}} = \widehat{\Phi}_{A\to T}^{\mathrm{T}}(X_A^{\mathrm{T}}), \tag{17}$$

$$Y^{\mathrm{T}} = \widehat{\Phi}_{T\to T+R/8}^{\mathrm{T}}(X_T^{\mathrm{T}}), \tag{18}$$

$$U^{\mathrm{T}} = \frac{Y^{\mathrm{T}} - X_T^{\mathrm{T}}}{R/8}. \tag{19}$$

RK4 uses FP32 and equal macrosteps no longer than $1/16$. The shared CIFAR field $f_\theta(t, x, R)$ adds the span to its 512-dimensional time-conditioning path through

$$z(R) = \frac{R - 2.5}{1.5}, \qquad e(R) = W_2\,\mathrm{SiLU}(W_1 z(R) + b_1) + b_2. \tag{20}$$

The widths $1 \to 32 \to 512$ add 16,960 parameters. Zero-initialized W_2, b_2 recover the teacher at every horizon before training.

A.2 Teacher preparation and adaptation cost

We separate teacher preparation, student adaptation, and generation. One teacher-example evaluation is one network evaluation on one image; batching changes invocation counts but not this quantity. Table 4 reports measured construction costs for the common teacher paths. The 64×64 caches reuse the first RK4 stage as the left-node teacher velocity. Their generation times sum the four shard timers and exclude setup and final serialization.

Table 4: **Shared teacher-path preparation.** Each cache contains 12,288 RK4-64 paths. Times are measured GPU-seconds; storage uses decimal GB.

Dataset	Teacher-example evaluations	GPU-seconds	Storage (GB)
CIFAR-10	786,432	259.9	1.359
FFHQ-64	786,432	2,552.1	10.268
AFHQv2-64	786,432	2,551.5	10.268

CIFAR-10 additionally integrates its continuous-start and continuous-horizon labels for each training sequence. Each example needs at most two RK4 steps from its cached anchor, eight from start to endpoint, and one teacher-velocity evaluation. The resulting upper bound is 41 teacher-example evaluations per label, or 7,872,000 for a 3,000-update sequence of batches of 64. The 128-state diagnostic set is prepared separately. Both targets reuse this label artifact in the paired study.

Table 5 counts student work separately from teacher preparation. Both objectives backpropagate through one prediction per example. CIFAR-10 uses forward-mode input differentiation; the 64×64 implementation uses reverse-over-reverse differentiation in microbatches of 16. The derivative traversals are additional work beyond the primal-forward counts.

Each original cache is shared across six runs; continuous CIFAR-10 labels are shared within a seed pair. The latter cost 3438.3 ± 35.7 GPU-seconds and 7.872 ± 0.000 million teacher-example evaluations per seed. A standalone run incurs its required preparation cost. Teacher pretraining, rate selection, held-out probe preparation, and generation are separate.

Table 5: **Student adaptation cost.** Training time and peak allocated GPU memory are means ± sample SD across three training seeds on A100 40GB GPUs. Each target receives 3,000 updates. Primal and input-JVP examples are in thousands; JVP calls include microbatching. Teacher preparation is reported in Table 4 and the accompanying text.

Dataset	Target	Time (s)	Peak (GiB)	Primal (k)	JVP (k)	JVP calls
CIFAR-10	Endpoint	276.7 ± 3.6	6.74 ± 0.00	192	0	0
	Stopped JVP	775.5 ± 9.2	6.74 ± 0.00	384	192	3,000
FFHQ-64	Endpoint	1349.2 ± 10.8	24.80 ± 0.00	96	0	0
	Stopped JVP	9902.6 ± 54.7	24.81 ± 0.00	192	96	6,000
AFHQv2-64	Endpoint	1329.8 ± 14.0	24.80 ± 0.00	96	0	0
	Stopped JVP	9860.6 ± 36.2	24.81 ± 0.00	192	96	6,000

Training time includes data transfer and earlier checkpoint/probe overhead, and excludes setup and the final save/probe. Each update has one loss backward. Primal counts exclude derivative traversals and are not FLOPs.

A.3 Supervision precision and measured learning time

Both comparisons use the same fixed CIFAR-10 teacher and 1,000 cached batch rows. Each row contains 32 interior states and 32 independent anchor states. Interior horizons follow $U[1/128, 1/2]$, with start times uniform on $[0, 1 - h]$. Anchor horizons equal $h_0 = 1/128$. The loss averages the interior term and the fixed endpoint-anchor term with equal weights. Endpoint training uses integrated labels in both terms; stopped-JVP training uses its differential target in the interior. All network parameters are trained in FP32, with TF32 disabled. Both targets use AdamW, 100 linear warmup updates followed by a constant rate, gradient clipping at 1, and no EMA. Precision runs use 10^{-5}; the main time-budget comparison selects a rate separately for each target.

For the precision comparison, each seed runs 1,000 updates on both label sets. Reference RK4 integration uses a maximum step of $1/128$ and at least eight substeps per interval. The alternative takes two RK4 substeps, or eight teacher evaluations per interval. Both models share the same initialization and batch order within each pair. Three pairs vary the initialization and order on this common dataset.

The time-budget comparison uses three further paired seeds after rate selection. Required interior labels are produced on first use of each batch row and cached on CPU. Endpoint pays eight teacher calls per new row; stopped JVP pays one velocity call. Later visits reuse the teacher quantities; the stopped-JVP target is recomputed from the current student at every update. Each epoch independently permutes all 1,000 rows; the two targets share the same sequence prefix, even when their update counts differ. Each pair trains sequentially on one A100 GPU, with evaluation beginning after all timed training.

Measured learning time includes label production, transfers, optimizer updates, and trace recording. Shared states and endpoint-anchor labels are prepared before timing begins. Initialization, input preload, checkpoint serialization, and evaluation lie outside this learning clock. Teacher-forward counts distinguish common data construction from interior-label production. Checkpoints at 300 and 600 seconds retain the last fully completed update inside each budget. Unused time is less than 1% in every run. Both checkpoints are evaluated with 50,000 common Gaussian inputs at NFE 5 and 7. The precision and time-budget comparisons each use their own fixed generation stream. Reported uncertainty is sample SD across training seeds on the common dataset.

Independent rate selection. Each target tests $\{2.5 \times 10^{-6}, 5 \times 10^{-6}, 10^{-5}, 2 \times 10^{-5}\}$ on one separate initialization and minibatch order, using 600 seconds per candidate. The mean of NFE-5 and NFE-7 FID on 2,000 common generated images selects the rate; exact ties favor the smaller rate. Both targets select 2×10^{-5}, the upper endpoint of the tested range (Table 6). Selection uses training seed 940101 and generation seed 941001. The three subsequent pairs use initialization/order seeds 940401, 940501 and 940601, and a separate 50,000-image generation stream (seed 941401). All runs use the shared state–anchor dataset from seed 930101. Within each pair, the targets share initialization and minibatch order. The pairs alternate which target trains first on the same A100. Table 7 gives each checkpoint's results.

Table 6: **Independent learning-rate selection.** Each candidate uses the same 600-second training protocol and 2,000 generated images at each NFE. For each target, the smallest mean of the two FIDs selects the rate; exact ties choose the smaller rate. These pilot models are separate from the three-seed comparison.

Target	LR	NFE 5	NFE 7	Mean	Selected
Endpoint	2.5×10^{-6}	23.690	21.678	22.684	
Endpoint	5×10^{-6}	21.882	20.611	21.247	
Endpoint	1×10^{-5}	20.729	19.750	20.240	
Endpoint	2×10^{-5}	20.298	19.488	19.893	Yes
Stopped JVP	2.5×10^{-6}	26.265	28.144	27.204	
Stopped JVP	5×10^{-6}	23.812	26.188	25.000	
Stopped JVP	1×10^{-5}	24.395	23.540	23.968	
Stopped JVP	2×10^{-5}	21.674	20.925	21.300	Yes

Table 7: **Individual equal-time results.** Each row is one model checkpoint; the two FIDs use 50,000 images each. Independent pilot tuning selects 2×10^{-5} for both targets. Label seconds are included in the training budget. Bold marks lower FID within each pair.

					Field	FID ↓	
Time (s)	Seed	Target	Updates	Label (s)	RMSE ↓	NFE 5	NFE 7
300	940401	Endpoint	635	127.0	0.0565	**8.068**	**6.912**
		Stopped JVP	683	17.9	0.0704	13.085	12.563
	940501	Endpoint	637	126.9	0.0604	**8.656**	**7.258**
		Stopped JVP	687	17.7	0.0654	11.887	11.682
	940601	Endpoint	638	126.8	0.0557	**8.220**	**7.085**
		Stopped JVP	688	17.8	0.0665	11.790	10.934
600	940401	Endpoint	1473	199.8	0.0474	**7.106**	**6.230**
		Stopped JVP	1390	26.1	0.0605	7.973	7.370
	940501	Endpoint	1480	199.2	0.0486	**7.143**	**6.318**
		Stopped JVP	1399	25.7	0.0603	8.678	7.972
	940601	Endpoint	1483	198.6	0.0478	**7.029**	**6.129**
		Stopped JVP	1399	25.8	0.0601	7.814	7.212

Sensitivity to learning rate. Tables 8 and 9 complement the independently tuned comparison with a shared rate and a half-rate JVP arm. The shared-rate study uses three paired seeds. Its student updates average 0.272 seconds for endpoint and 0.418 for JVP; the clock also includes labels and transfers. The half-rate study uses one pair and the same 600-second budget, with 50,000 common generation inputs. Halving the JVP rate improves NFE-5 FID while slightly worsening NFE-7 FID.

Table 8: **Equal-time comparison at a common learning rate.** Both targets use 10^{-5} on three paired seeds. FID uses 50,000 images; uncertainty is sample SD across training seeds. Updates and included label times are means.

Budget			Labels	Field	FID ↓	
(s)	Target	Updates	(s, included)	RMSE ↓	NFE 5	NFE 7
300	Endpoint	627.7	127.5	0.0748 ± 0.0031	$\mathbf{9.854 \pm 0.579}$	$\mathbf{8.146 \pm 0.415}$
	Stopped JVP	670.7	18.2	0.0752 ± 0.0009	10.319 ± 1.914	12.951 ± 1.751
600	Endpoint	1445.3	203.1	0.0554 ± 0.0023	$\mathbf{8.249 \pm 0.649}$	$\mathbf{7.148 \pm 0.475}$
	Stopped JVP	1364.3	26.9	0.0661 ± 0.0003	12.680 ± 0.918	12.272 ± 0.962

Table 9: Learning-rate comparison at a 600-second budget.

Target	Learning rate	NFE-5 FID ↓	NFE-7 FID ↓
Endpoint	10^{-5}	**7.581**	**6.648**
Stopped JVP	10^{-5}	11.743	11.254
Stopped JVP	5×10^{-6}	8.928	11.651

B Finite-Model and Target Analysis

B.1 Attainable teachers and finite-model equilibria

Polynomial control. The student has three output coordinates, each a polynomial of degree three in $z = 2h$, with $h \in [1/8, 1/2]$. We orthonormalize the monomial design on 64 equally spaced training horizons. The teacher belongs to the same 12-parameter class. Both targets receive the identical teacher array u and use $\mathcal{L} = -h\partial_h$, so $\mathcal{Q}_\lambda(g) = u + \lambda\mathcal{L}(g - u)$. The operator $I - \lambda\mathcal{L}$ has eigenvalues $1 + \lambda k > 0$; therefore the teacher is its unique fixed-point function for every tested λ. Full-batch gradient descent on the half-squared stopped loss is exactly the coefficient relaxation because the design is orthonormal.

Five initialization seeds, $104729, 130363, 155921, 181081, 206369$, run for 800 updates. The tested learning rates are $0.10, 0.20, 0.35, 0.45, 0.49, 0.51, 0.60, 0.79, 0.81, 1.00, 1.50, 1.90, 2.01$. At $\alpha = 0.20$, both endpoint and stopped objectives reach the same coefficient solution within 2.65×10^{-15} across all seeds. The intervals $[1.90, 2.01]$, $[0.79, 0.81]$, and $[0.49, 0.51]$ bracket the stable-to-unstable transition for $\lambda = 0, 0.5, 1$ respectively.

Nonlinear neural control. A width-four tanh network maps two characteristic labels ξ to four coefficients of a cubic horizon polynomial. All hidden and output weights are trained, giving 32 parameters. The teacher uses the same architecture with fixed parameters. A normalized Legendre horizon basis conditions the polynomial evaluation; the neural Jacobian is neither whitened nor fixed. Both arms receive the same $u(\xi, h)$ and $\partial_h u(\xi, h)$, and regress to $u - \lambda h\partial_h(g - u)$ with the complete label detached. For each ξ, the equation $e + \lambda h\partial_h e = 0$ has only the zero cubic solution. The teacher is the unique fixed-point function in this class, although its network parameters can have equivalent representations.

The labels ξ are invariant along a teacher characteristic. One explicit flow realization takes state (ξ, y), terminal time s, and $y(t) = -(s - t)u(\xi, s - t)$. Its velocity in the y direction is $u + h\partial_h u$, with zero velocity in ξ. The characteristic derivative of a predictor $g(\xi, h)$ is therefore $-\partial_h g$.

Training uses an 8×8 Gauss–Legendre grid for ξ and eight horizon nodes on $(0, 1)$. Evaluation uses an independent 11×11 midpoint grid and 31 horizon nodes. Each initialization seed perturbs the teacher parameters with independent $\mathcal{N}(0, 0.08^2)$ noise. Both arms use Adam (Kingma & Ba, 2015) with $\beta = (0.9, 0.999)$, $\epsilon = 10^{-8}$, no weight decay, and the same cosine learning-rate schedule from 0.003 to 3×10^{-6} over 30,000 updates. This is a local convergence experiment around an attainable teacher.

Table 10: **Learning toward an attainable neural teacher.** Held-out RMS divided by the teacher RMS, mean ± sample SD over five parameter-initialization seeds.

Adam updates	Fixed endpoint	Stopped differential target
0	0.42344 ± 0.14493	0.42344 ± 0.14493
100	0.02500 ± 0.00465	0.08250 ± 0.03153
1,000	0.004398 ± 0.001777	0.018364 ± 0.004925
5,000	0.000185 ± 0.000126	0.004531 ± 0.003757
10,000	0.000356 ± 0.000322	0.001652 ± 0.001702
30,000	$(8.51 \pm 2.08) \times 10^{-15}$	$(1.76 \pm 2.11) \times 10^{-5}$

The largest final relative RMS across all ten runs is 5.19×10^{-5}. The weighted function kernel $W^{1/2}JJ^\top W^{1/2}$ changes by 7.27–10.61% in Frobenius norm. Independent finite differences check the stopped gradient, forward Jacobian, and horizon derivative with errors below 6×10^{-10}. Injecting

a pure h^k error through the output bias gives gradient ratios $2, 3, 4$. Actual SGD function-update ratios are $2.000011, 3.000020, 4.000021$; the first Adam step gives $1.000006, 1.000008, 1.000014$ because its coordinate normalization removes almost all uniform gradient rescaling.

Transported approximation error. Let the fitted class be $f_0 + V$, and let P denote orthogonal projection onto the finite trainable subspace V. The endpoint regression solution f_{end} has residual $r = f_{\text{end}} - u$ with $Pr = 0$. Writing the stopped stationary function as $f_{\text{end}} + \delta$, its semi-gradient condition is $P(I - \lambda\mathcal{L})(r + \delta) = 0$, or

$$(I_V - \lambda P\mathcal{L}|_V)\,\delta = \lambda P\mathcal{L}r. \tag{21}$$

If the left-hand map is invertible, this determines the equilibrium shift. Exact realizability sets $r = 0$ and $\delta = 0$. Otherwise, the target can return an unrepresentable residual to a trainable direction. A separate quartic-teacher/cubic-student control verifies this formula: actual stopped gradient descent reaches the predicted stationary functions, with maximum dense-grid shift error 7.11×10^{-15}.

B.2 Relation to finite-time learning objectives

Table 11 uses increasing-time coordinates and a fixed teacher direction. $D\mathcal{Q}$ differentiates label recomputation between updates; the label remains detached within each update.

Table 11: **Target operators in common coordinates.** Here $u = u_h^{\mathrm{T}}$, $v = v_{\mathrm{T}}$, and $P_a g = g(t + ah, x + ahv, (1-a)h)$. The last column is the diagonal order-k gain of the stopped regression residual, including the stated scalar loss weight. It precedes model and optimizer preconditioning.

Construction	Recomputed label	Label derivative	Residual gain
Integrated teacher	u	0	1
Differential	$v + h\mathcal{D}_s g$	$h\mathcal{D}_s$	$1 + k$
Anchored feedback	$u + \lambda h\mathcal{D}_s(g - u)$	$\lambda h\mathcal{D}_s$	$1 + \lambda k$
Finite shift, weight $1/a$	$av + (1-a)P_a g$	$(1-a)P_a$	$\sum_{j=0}^{k}(1-a)^j$

The rows cover integrated teacher labels in CMT and Int-MeanFlow (Hu et al., 2026; Wang et al., 2026, Eq. (8)), MeanFlow's differential identity (Geng et al., 2025), and the fixed-direction version of AlphaFlow's Eq. (8) (Zhang et al., 2026). Appendix B.3 derives the finite-shift forcing relative to u.

Learning the derivative direction. The instantaneous-velocity loss rewrite in iMF retains the original semi-gradient for a fixed direction. Its full objective also learns that direction (Geng et al., 2026, Eqs. (9)–(12)). Write the direction as $b = v + \delta b$ and the field as $g = u + e$. Then

$$\begin{aligned} \mathcal{Q}(g, b) - u &= h\mathcal{D}_s e + h\,\delta b \cdot \nabla_x u + h\,\delta b \cdot \nabla_x e, \\ D\mathcal{Q}[(u, v)](e, \delta b) &= h\mathcal{D}_s e + h\,\delta b \cdot \nabla_x u. \end{aligned} \tag{22}$$

A learned direction introduces a coupled velocity-error term in addition to horizon feedback. If $b = g(t, x, 0)$ and the boundary is anchored to v, this term vanishes. Otherwise it participates in the joint learning dynamics. Our paired image experiments freeze b to isolate the first term.

B.3 Finite consistency shifts: response and forcing

We analyze the fixed-direction case of the finite consistency target in AlphaFlow, Eq. (8) and Algorithm 1 (Zhang et al., 2026). In increasing-time coordinates, let h denote the signed interval length and $a \in (0, 1]$ its shift ratio. Define

$$(P_a g)(t, x, h) = g(t + ah, x + ahv(t, x), (1-a)h), \tag{23}$$

$$\mathcal{Q}_a g = av + (1-a)P_a g, \qquad \mathcal{R}_a(g) = \frac{g - \mathcal{Q}_a g}{a}. \tag{24}$$

Here v is a smooth teacher field independent of the student. For a decreasing-time interval, $h < 0$ gives the backward Euler shift with the same formulas. For the half-squared loss weighted by $1/a$, stopping the label gives semi-gradient $J_\theta^* \mathcal{R}_a(g_\theta)$. Label recomputation across updates has derivative $D\mathcal{Q}_a = (1-a)P_a$, even though this derivative is blocked within each gradient step.

Horizon-order response. Write a degree-K error jet as $e(t,x,h) = \sum_{k=0}^{K} h^k e_k(t,x) + O(h^{K+1})$, with sufficiently smooth coefficient fields. For an isolated coefficient,

$$(1-a)P_a(h^k e_k) = (1-a)^{k+1} h^k \left[e_k + ah(\partial_t + v \cdot \nabla_x) e_k + O(h^2) \right]. \tag{25}$$

All terms generated by the state/time displacement have degree greater than k. Therefore the jet map is triangular, with diagonal blocks $(1-a)^{k+1} I$. For $\mathcal{H}_a = D\mathcal{R}_a = [I - (1-a)P_a]/a$, these become

$$\beta_k(a) I, \qquad \beta_k(a) = \frac{1-(1-a)^{k+1}}{a} = \sum_{j=0}^{k} (1-a)^j. \tag{26}$$

For $k \geq 1$, differentiating the finite sum shows that β_k is strictly decreasing on $(0,1)$, from $k+1$ to one. The degree-zero gain is one throughout. Coefficient coupling remains below the diagonal; Eq. (26) describes the diagonal response, not the complete evolution of a general superposition of errors.

At fixed nonzero h, Taylor expansion in a gives $P_a g = g + ah\mathcal{D}_s g + O(a^2)$, where $\mathcal{D}_s = \partial_t + v \cdot \nabla_x - \partial_h$. Substitution gives

$$\mathcal{Q}_a g = g + a(v - g + h\mathcal{D}_s g) + O(a^2), \qquad \mathcal{R}_a(g) \longrightarrow g - v - h\mathcal{D}_s g. \tag{27}$$

The known differential limit is therefore resolved into separate horizon-order gains by Eq. (26).

The integrated teacher and the finite-shift forcing. Let $u(t,x,h) = [\Phi(t,t+h,x) - x]/h$ be the exact teacher average and define $b_a = \mathcal{Q}_a u - u$. Affinity of the target gives the exact decomposition

$$\mathcal{R}_a(u+e) = \mathcal{H}_a e - \frac{b_a}{a}. \tag{28}$$

Put $A = \partial_t + v \cdot \nabla_x$ and $c = Av$. The teacher-flow expansion and the Euler state shift imply

$$u(t,x,h) = v + \tfrac{h}{2} c + O(h^2), \tag{29}$$

$$v(t+ah, x+ahv) = v + ahc + O(h^2), \tag{30}$$

$$u(t+ah, x+ahv, (1-a)h) = v + \tfrac{1+a}{2} hc + O(h^2). \tag{31}$$

Multiplying the last expression by $1-a$, adding av, and subtracting the first gives

$$b_a = -\frac{a^2 h}{2} Av + O(h^2) \qquad \text{at fixed } a. \tag{32}$$

Also $b_0 = 0$, and the exact teacher identity $u = v + h\mathcal{D}_s u$ gives $\partial_a b_a|_{a=0} = 0$. This gives $b_a = O(a^2)$ for fixed h, and the normalized forcing vanishes as $a \to 0$. At $a = 1$, $\mathcal{Q}_1 g = v$ and $b_1 = v - u$, which is generally nonzero. The unit gain at that endpoint is accompanied by a different target field.

The forcing has a direct geometric source: the Euler prefix starts the remaining teacher trajectory at $x + ahv$, rather than at $\Phi(t, t+ah, x)$. Replacing that prefix by the exact teacher flow and using its average velocity preserves u by the flow composition law. The diagonal gains $\beta_k(a)$ remain the same, while b_a vanishes. The two terms in Eq. (28) separate label response from teacher-integration error.

Measurement with a nonnormal affine flow. We use the exact flow of $x' = Mx + d$ in two dimensions, where

$$M = \begin{pmatrix} -.4 & 1.1 \\ -.7 & -.2 \end{pmatrix}, \qquad d = \begin{pmatrix} .3 \\ -.15 \end{pmatrix}.$$

An augmented matrix exponential gives $\Phi_h(x)$ and the teacher average. Eight states are sampled uniformly from $[-.8, .8]^2$ with seed 104729. The injected errors are $e_k(x,h) = h^k p(x)$ for $k = 1,2,3$, with $p(x) = (1 + .25\sin x_0 + .2x_1,\ -.4 + .3\cos x_1 - .15x_0)^\top$. For each $a \in \{.05, .1, .2, .4, .6, .8, 1\}$, we independently evaluate $\mathcal{R}_a(u+e_k)$ and $\mathcal{R}_a(u-e_k)$. Their half-difference removes b_a exactly. Dividing by h^k and fitting a cubic in $h/.08$ on $h \in \{.08, .04, .02, .01, .005, .0025\}$ estimates the leading vector coefficient. Projecting it onto $p(x)$ gives the measured gain in Figure 4a. The largest gain discrepancy is 4.30×10^{-8} across the $7 \times 3 \times 8$ measurements; the largest relative vector discrepancy is 1.66×10^{-8}.

Forcing is evaluated separately on six halving horizons from .1 to .003125. Figure 4b plots $\|b_a\|/\|Mv\|$, preserving all eight statewise curves. A cubic fit of $b_a/(a^2h)$ recovers the coefficient $-Mv/2$ with maximum relative error 1.55×10^{-8}. The fitted small-horizon exponents range from 0.9966 to 0.9998; the step-ratio exponents range from 2.00005 to 2.00033. An independent centered directional derivative at $h = .2$ confirms Eq. (27) with median convergence rate $a^{0.9989}$. These measurements isolate the refreshed-label operator; the full AlphaFlow training system also applies its chosen shift direction, adaptive weighting, network Jacobian, and optimizer to this response.

C Image-Target Diagnostics and Comparisons

Table 12: **Paired finite-step flow learning across domains and sampling budgets.** FID after 3,000 updates; CIFAR-10 NFE 7 is the secondary budget. Each seed averages two common streams of 50,000 images on CIFAR-10 and 10,000 on the 64×64 domains. Summaries give mean ± sample SD and paired 95% Student-t intervals across three training seeds. Positive gaps favor endpoint.

			FID ↓		
Dataset	NFE	Pair	Fixed endpoint	Stopped JVP	Gap / 95% CI
CIFAR-10	5	1	**7.665**	24.142	+16.48
		2	**7.625**	18.184	+10.56
		3	**7.480**	21.072	+13.59
		Mean ± SD	**7.590 ± 0.098**	21.133 ± 2.979	+13.54 [+6.19, +20.89]
CIFAR-10	7	1	**7.819**	31.069	+23.25
		2	**7.760**	22.140	+14.38
		3	**7.645**	27.176	+19.53
		Mean ± SD	**7.741 ± 0.089**	26.795 ± 4.476	+19.05 [+7.99, +30.12]
FFHQ-64	8	1	**5.586**	33.801	+28.22
		2	**5.543**	33.446	+27.90
		3	**5.521**	33.931	+28.41
		Mean ± SD	**5.550 ± 0.033**	33.726 ± 0.251	+28.18 [+27.54, +28.81]
AFHQv2-64	8	1	**3.988**	14.992	+11.00
		2	**3.882**	14.553	+10.67
		3	**3.975**	15.216	+11.24
		Mean ± SD	**3.948 ± 0.058**	14.921 ± 0.337	+10.97 [+10.26, +11.68]

C.1 Continuous interval support

We also train the FFHQ-64 pair with continuously sampled interval starts and ends. The architecture, optimizer, and initialization follow Appendix A.1. Both arms receive 1,000 updates with batch size 32 and FP32 arithmetic. They share the training sequence and the 2×10^{-5} peak learning rate. We use the first 1,000 updates of the existing schedule: 100 warmup updates followed by the 3,000-update cosine decay.

For each cached interval $[\sigma_i, \sigma_{i+1}]$, draw independent fractions $a \sim U(0, 0.2)$ and $b \sim U(0.8, 1)$. With $\delta_i = \sigma_{i+1} - \sigma_i$, the start is $\sigma_i + a\delta_i$ and the end is $\sigma_i + b\delta_i$. The signed horizon is $(b - a)\delta_i$ in the EDM noise coordinate. RK4 first transports the cached state to the sampled start, then integrates from start to end. Both objectives reuse these states, the endpoint displacements, and the teacher velocities at the starts.

We choose the integration resolution on 32 held-out states. Doubling RK4 from two to four microsteps per segment changes the endpoint field by 3.573×10^{-4} RMSE. This is 1.001% of the 3.568×10^{-2} separation between the two 3,000-update FFHQ reference fields. This is below the chosen 5% tolerance; preparation uses two microsteps per segment and takes 1,797.6 GPU-seconds and 544,000 teacher-example evaluations.

Endpoint training retains the lower field error and FID when starts and horizons vary continuously. Stopped JVP attains the lower differential residual, preserving the distinction between these two error measures.

Table 13: **Held-out field error and differential residual.** Each metric is a root mean squared error over the same 128 held-out states, shown in units of 10^{-3}. Entries are means $\pm$ sample SD across three independent training seeds. Field error measures approximation to the numerical teacher endpoint field; differential residual measures satisfaction of its transport identity.

		Field error ↓		Differential residual ↓	
Dataset	Updates	Endpoint	JVP	Endpoint	JVP
CIFAR-10	1k	51.97 ± 0.43	145.43 ± 2.06	210.72 ± 1.18	511.86 ± 13.30
	3k	50.08 ± 0.29	132.75 ± 3.96	198.49 ± 1.06	401.71 ± 4.56
FFHQ-64	1k	12.31 ± 0.07	36.11 ± 0.28	30.76 ± 0.14	22.28 ± 0.25
	3k	11.55 ± 0.01	34.01 ± 0.84	30.16 ± 0.20	21.31 ± 0.12
AFHQv2-64	1k	12.76 ± 0.04	28.76 ± 1.02	29.75 ± 0.26	23.01 ± 0.22
	3k	11.83 ± 0.02	27.43 ± 0.65	28.94 ± 0.18	22.59 ± 0.71

Table 14: **Continuous-start and continuous-horizon FFHQ-64 training.** One paired training seed, 1,000 updates, and NFE 8. FID uses the same 2,000-image generation stream for both arms. Field error and differential residual use the same 128 held-out states and are shown in units of 10^{-3}.

FID ↓		Field error ↓		Differential residual ↓	
Endpoint	Stopped JVP	Endpoint	Stopped JVP	Endpoint	Stopped JVP
12.572	24.680	**15.985**	31.317	35.518	**21.864**

C.2 Boundary-connected target comparison

Training on connected intervals. One CIFAR-10 pair uses the connected-interval sampler and equally weighted interior/anchor loss in Appendix A.3, with $h_0 = 1/128$. Both arms share initialization and inputs, and train for 1,000 FP32 AdamW updates at 10^{-5} after 100 warmup updates. The same 256 interior and 256 boundary states are evaluated at updates 200 and 1,000. Stopped JVP initially has lower interior error; endpoint overtakes it by update 1,000.

	Interior MSE ↓ (10^{-3})		Boundary MSE ↓ (10^{-3})	
Updates	Endpoint	Stopped JVP	Endpoint	Stopped JVP
200	10.333	**8.674**	**2.705**	12.225
1,000	**3.758**	4.753	**1.294**	7.003

Boundary transport at 1,000 updates. The learned boundary RMSEs are 0.03597 and 0.08368, respectively. We measure how their boundary difference propagates into the interior.

Let g_E, g_J denote the trained fields and write $t' = t + h - h_0$. Using the frozen teacher flow, define

$$b(t, x, h) = \frac{h_0}{h}(g_J - g_E)\big(t', \Phi^{\mathrm{T}}_{t \to t'}(x), h_0\big), \qquad g_J^{\mathrm{cf}} = g_J - b. \tag{33}$$

For smooth fields, $b - h\mathcal{D}_s b = 0$, where $\mathcal{D}_s = v_{\mathrm{T}} \cdot \nabla_x + \partial_t - \partial_h$. The replacement matches g_J^{cf} to g_E at h_0 while preserving the differential residual $r(g) = g - v_{\mathrm{T}} - h\mathcal{D}_s g$.

Let $M(g)$ be mean squared pixelwise error against u_h^{T} on the same held-out states. Define the original gap $\Delta = M(g_J) - M(g_E)$ and its reduction $C = M(g_J) - M(g_J^{\mathrm{cf}})$. The table measures this direct boundary contribution on all states and on the $h \geq 1/8$ subset.

States	n	$10^3\Delta$	$10^3 C$	C/Δ
All	256	0.994928	0.152260	15.30%
$h \geq 1/8$	192	0.414050	0.019271	4.65%

C.3 Injected-mode response measurements

Image-model response measurement. At the 1,000- and 3,000-update endpoint checkpoints, we add a parameter-independent output perturbation $\delta_k(h) = \epsilon h^k u / \sqrt{\mathbb{E}_h[h^{2k}]}$, where $h \sim U[1/8, 1/2]$,

$k \in \{1, 2, 3\}$, and u is spatially constant with RGB values $(1, -1, 1/2)$ normalized to unit RMS. For each $\lambda \in \{0, 1/2, 1\}$, we compare one update with this overlay to an unperturbed update on the same minibatch. Each branch restores the checkpoint weights, buffers, AdamW moments, and step counter. AdamW uses learning rate 2×10^{-6}, $\beta = (0.9, 0.999)$, numerical $\epsilon_{\mathrm{Adam}} = 10^{-8}$, weight decay 0.01, and gradient-norm clipping at 1. Training uses bfloat16 autocast with float32 MSE; post-update outputs use float32 and inner products use float64.

Let d_λ be the perturbed-minus-unperturbed parameter gradient, measured before clipping. Let $g^+_{\lambda,\delta}$ and $g^+_{\lambda,0}$ denote the two post-update fields, with the overlay retained in $g^+_{\lambda,\delta}$. On a separate held-out batch, define

$$G_\lambda = \frac{\langle d_\lambda, d_0 \rangle}{\|d_0\|^2}, \qquad R_\lambda = \frac{1 - m_\lambda}{1 - m_0},$$
$$m_\lambda = \frac{\langle g^+_{\lambda,\delta} - g^+_{\lambda,0}, \delta_k \rangle}{\|\delta_k\|^2}.$$

Here G_λ measures gradient amplification; R_λ compares the overlay's removal after the actual update with its removal under fixed endpoint supervision. The probe uses four update/held-out batch pairs of 16 examples and amplitudes $\epsilon \in \{0.01, 0.03\}$. Figure 5b reports G_1 and R_1 at $\epsilon = 0.03$: medians and interquartile ranges pool four batches at each of two checkpoints. Clipping is inactive in these trained-checkpoint measurements.

Table 15: **A state-dependent mode recovers the degree-raising term.** We inject $\delta g = 0.05 z \phi(x, t)$ on 128 held-out initialization states, with $\phi(x, t)$ given by the student's normalized zero-horizon field. An independent neural JVP measures $\mathcal{A}\phi = (\partial_t + v_{\mathrm{T}} \cdot \nabla_x)\phi$; a degree-4 fit separates the degree-1 diagonal from the degree-2 material derivative. Intervals resample states 2,000 times.

Response component	Degree	Gain (95% interval)	Cosine	RMS error
Diagonal $-z\phi$	1	0.99999987 [0.99999983, 0.99999991]	1.00000000	6.96×10^{-6}
Raising $(1/8)z^2\mathcal{A}\phi$	2	1.00000056 [1.00000039, 1.00000075]	1.00000000	2.12×10^{-5}

C.4 Natural errors and actual updates

Fixed-state first-order coefficients. At 512 shared diagnostic states, we fit degree-four polynomials in $z = 8h$ to the natural error $e = g - u$ and label error $q = \mathcal{T}g - u$. Each (t, x) stays fixed. Nodes are $h_j = [1 - \cos(j\pi/10)]/16$, $j = 0, \ldots, 10$, with the exact $h = 0$ teacher boundary. For the linear coefficients e_1, q_1, define

$$\rho_1 = \frac{\|e_1\|^2}{\|e_1\|^2 + \|q_1 + e_1\|^2},$$

where norms sum over states and output coordinates. Endpoint checkpoints at 1,000 and 3,000 updates give 25.12% and 24.20%; the paired JVP checkpoints give 18.52% and 18.96%. The denominator sums the squared norms of two generally nonorthogonal parts of the same first-order coefficient. The remainder $q_1 + e_1$ includes material transport and finite-window fitting error. This small-horizon measurement covers $0 \le h \le 1/8$.

Transport along teacher characteristics. For a teacher characteristic with fixed right endpoint s, let $\bar{e}(h) = e(s - h, X_{s-h}, h)$. The chain rule gives

$$\frac{\mathrm{d}\bar{e}}{\mathrm{d}h} = \partial_h e - (\partial_t + v_{\mathrm{T}} \cdot \nabla_x)e = -\mathcal{D}_s e, \qquad \bar{r}(h) = \bar{e}(h) + h\frac{\mathrm{d}\bar{e}}{\mathrm{d}h}.$$

This derivative includes state and time variation, unlike the fixed-state jet above. Nine calibration horizons between $1/8$ and $1/2$ supply endpoint errors; eight interleaved horizons test the reconstructed residual. A cubic fit and a piecewise-linear interpolant use the same calibration errors. The four checkpoints come from one training seed and share 16 held-out paths. Eight paths supply update gradients; the other eight measure post-update field error.

Let J be the student Jacobian on the update states. We compare $J^\top \widehat{r}$ with the exact stopped gradient $J^\top r$. The endpoint comparator uses $J^\top e$ on those same states. Residual errors use all 16 paths; gradient cosines use the eight update paths. Network evaluations and optimizer steps use FP32 with TF32 disabled.

		Residual relative L^2 error ↓		Cosine with stopped gradient ↑	
Trained target	Updates	Cubic	Linear	Cubic	Endpoint
Endpoint	1,000	0.404	0.167	0.9927	0.6797
Endpoint	3,000	0.433	0.184	0.9940	0.6856
Stopped JVP	1,000	0.599	0.255	0.9345	0.4211
Stopped JVP	3,000	0.528	0.206	0.9829	−0.0426

Linear interpolation lowers residual error; the cubic fit closely matches the gradient direction. All one-step comparisons in this subsection restore checkpoint weights, buffers, AdamW moments, and step count before each branch. At the 1,000-update endpoint checkpoint with rate 8×10^{-6}, exact, cubic, and scalar-rescaled endpoint residuals reduce held-out MSE by 3.3018, 3.2951, and 4.4668 times 10^{-5}. For training-choice comparisons, update paths select among $\lambda \in \{0, 1/2, 1\}$ and rates $\{5 \times 10^{-7}, 2 \times 10^{-6}, 8 \times 10^{-6}\}$; the other paths evaluate the chosen exact operation. Cubic and scalar predictors select identically at three checkpoints. At the fourth, cubic matches independent rate selection for endpoint regression.

Predicting a change in learning direction. We measure natural errors at 1,000 and 3,000 updates for three paired training seeds, giving twelve checkpoints. The 32 teacher paths are excluded from training. Sixteen paths supply updates; the other sixteen measure the resulting field error. Each set contains eight test horizons per path. Nine interleaved calibration horizons supply the cubic error fit. The same paths, horizons, and three learning rates are used for every checkpoint.

Write $G_e = \nabla_\theta F_U$, where F_U is endpoint field MSE on the update states. The fitted characteristic error and its transported residual define

$$\widehat{e}(h) = \sum_{k=0}^{3} a_k h^k, \qquad \widehat{r}(h) = \sum_{k=0}^{3} (k+1) a_k h^k, \qquad G_{\widehat{e}} = \frac{2}{m} J_U^\top \widehat{e}, \quad G_{\widehat{r}} = \frac{2}{m} J_U^\top \widehat{r}.$$

Here m counts output scalars. The exact stopped gradient is $G_r = 2J_U^\top r/m$. Since $\mathrm{d}F_U(\theta - \eta G_r)/\mathrm{d}\eta|_{\eta=0} = -\langle G_e, G_r \rangle$, a negative inner product identifies a local ascent direction for field error. Positive scalar rescaling preserves this sign.

Before evaluating exact test-horizon JVPs or evaluation-path student outputs, we save the cubic gradient prediction and its hypothetical AdamW steps on the update paths. The unweighted fit uses the same coefficients and Jacobian. At 3,000 updates, two JVP models have negative field-error projections, whereas the third remains positive. Cubic transport predicts all three signs; the unweighted fits remain nearly parallel to endpoint regression. The remaining checkpoints also have correctly predicted positive projections (Table 16).

Table 16: **Characteristic error transport predicts seed-dependent gradient directions.** Cosines with the endpoint field-error gradient on sixteen update paths. Unweighted and transported columns use the same cubic fit; Exact uses test-horizon JVPs.

Seed	Trained target	Updates	Unweighted fit	Transported fit	Exact stopped
920101	Endpoint	1,000	0.9982	0.6533	0.6753
920101	Endpoint	3,000	0.9985	0.7630	0.7651
920101	JVP	1,000	0.9975	0.1936	0.1501
920101	JVP	3,000	0.9997	−0.4980	−0.5255
920103	Endpoint	1,000	0.9980	0.6609	0.6824
920103	Endpoint	3,000	0.9985	0.7632	0.7679
920103	JVP	1,000	0.9973	0.1625	0.1549
920103	JVP	3,000	0.9993	0.2377	0.2278
920107	Endpoint	1,000	0.9978	0.7033	0.7232
920107	Endpoint	3,000	0.9984	0.7669	0.7692
920107	JVP	1,000	0.9974	0.1523	0.0925
920107	JVP	3,000	0.9996	−0.4175	−0.4570

A zero-gradient branch retains momentum and weight decay. Table 17 separates this baseline change from each target's contribution at the three late JVP checkpoints and three rates. For the two reversed

JVP models, exact and reconstructed transport both increase evaluation error relative to the zero-gradient branch at all three rates. Their raw changes are also positive. A scalar fit ce, using exact residuals at calibration horizons, gives the opposite relative effect. All fitted scalar coefficients are positive (1.106–3.694). This scalar changes the residual before AdamW; the three learning rates separately scale the restored optimizer step.

Table 17: **Finite AdamW responses at the three 3,000-update JVP checkpoints.** Field-MSE changes in units of 10^{-6}. Z_E is the raw zero-gradient change on evaluation paths. P_U is the cubic-minus-zero effect on update paths, saved before evaluation. E_{exact}, E_{cubic}, and E_{scalar} are observed target-minus-zero effects on sixteen separate evaluation paths. All learning rates were fixed before evaluation.

Seed	LR (10^{-6})	Z_E	P_U	E_{exact}	E_{cubic}	E_{scalar}
920101	0.5	8.736	0.534	2.447	2.314	−2.892
	2.0	35.103	2.244	9.900	9.364	−11.607
	8.0	142.986	10.634	41.247	39.000	−47.323
920103	0.5	−0.993	−2.289	−0.018	−0.001	−1.084
	2.0	−3.919	−9.111	−0.055	0.026	−4.346
	8.0	−14.811	−35.883	0.195	0.492	−17.420
920107	0.5	4.963	−0.106	1.748	1.638	−2.007
	2.0	19.950	−0.357	7.077	6.640	−8.049
	8.0	81.556	−0.090	29.573	27.728	−32.705

The evaluation effect matches the exact sign in 35/36 comparisons, whereas the saved update-path effect predicts 32/36 evaluation signs. The errors arise from reconstruction and transfer across paths. At seed 920107, cubic predicts negative update-path effects at all three rates, while exact effects are 0.004, 0.088, and 1.812 (10^{-6}): a reconstruction error. At seed 920103 and $\eta = 8 \times 10^{-6}$, cubic and exact effects are -35.883 and -37.390 on update paths, but exact evaluation gives $+0.195$: a transfer error. The remaining cubic–exact evaluation mismatch occurs near zero at seed 920103 and $\eta = 2 \times 10^{-6}$: $+2.575 \times 10^{-8}$ versus -5.545×10^{-8}.

On the earlier sixteen-path probe, seed 920101's 1,000-update JVP checkpoint at rate 8×10^{-6} gives exact/cubic effects $+1.666/-4.352$ (10^{-6}); both decrease error on the 32-path probe. Its late-checkpoint gradient reversal appears on both path sets. Transport can rotate gradients away from descent, while predicting small finite AdamW effects also requires accuracy along the evaluation-error direction.

Original minibatches and numerical precision. On the twelve checkpoints and original minibatches specified below, 128 held-out states evaluate each restored AdamW step in FP32 and BF16, with TF32 disabled. JVP has higher field error than endpoint in 23 of 24 FP32 comparisons and 16 of 24 BF16 comparisons. Repeated evaluations are identical; nine ordering changes arise from the forward precision. The late JVP model at seed 920101 also has endpoint–JVP gradient cosines -0.240 and $+0.231$ on its two batches. Both sampling and numerical precision affect these small update contrasts.

Order-aware reconstruction on training minibatches. The twelve checkpoints span three seeds, two targets, and 1,000 or 3,000 updates. Each uses the fifth and sixth training minibatches: 64 samples, no added boundary loss, $h \sim U[1/8, 1/2]$, and $t \mid h \sim U[0, 1-h]$. For each sample (t, x, h), fix $s = t + h$ and its cached endpoint $Y = x + hu_{\text{cache}}$. Nine Chebyshev–Lobatto horizons lie in $[1/8, \min(1/2, s)]$. Sixteen signed RK4 steps take x from t to each $t_j = s - h_j$, giving states x_j and errors $g_\theta(t_j, x_j, h_j) - (Y - x_j)/h_j$. A cubic least-squares fit in a normalized horizon coordinate gives $\widehat{e}$ and $\widehat{r} = \widehat{e} + h\,\mathrm{d}\widehat{e}/\mathrm{d}h$ at the original horizon. These coefficients include material variation along the characteristic; the formal jet holds (t, x) fixed.

Let J_U denote the whole-batch original-point BF16 student Jacobian and m the number of output scalars. The transported and unweighted gradients and their same-fit scalar comparator are

$$G_{\widehat{r}} = \frac{2}{m} J_U^\top \widehat{r}, \qquad G_{\widehat{e}} = \frac{2}{m} J_U^\top \widehat{e}, \qquad G_c = cG_{\widehat{e}}, \quad c = \max\left(0, \frac{\langle \widehat{e}, \widehat{r} \rangle}{\|\widehat{e}\|^2}\right).$$

The scale c uses all batch outputs. Fits use FP32 forwards and the original BF16 Jacobian. Predictions are saved before the exact stopped gradient G_r and post-update evaluation; the endpoint gradient G_e uses cached labels.

Sensitivity to teacher integration. Replaying teacher velocities and endpoints at the original RK4 maximum step $1/16$ reproduces every cached batch exactly. FP32 forwards allow cuDNN TF32; CUDA matrix multiplication TF32 is disabled. A fixed endpoint difference $\delta = Y_{16} - Y$ contributes δ/h to the cached-endpoint error relative to the sixteen-step endpoint. Although $I + h\partial_h$ annihilates this term, its cubic projection need not vanish. The RMS of $\delta(I + h\partial_h)P_3(1/h)$ is at most 1.39% of original field-error RMS, below the chosen 2% tolerance. Subtracting this endpoint-sensitivity gradient changes none of the 24 predicted signs. Optimizer branches use the original fits.

We report gradient reconstruction error and balanced accuracy of the projection sign separately, averaging over eight measurements per seed. The best nonnegative scaling after revealing G_r is $c_* = \max(0, \langle G_{\widehat{e}}, G_r\rangle / \|G_{\widehat{e}}\|^2)$. This gives the minimum error within the scalar family.

Table 18: **Gradient reconstruction beyond scalar rescaling.** Relative error is $\|G - G_r\|/\|G_r\|$, with the native BF16 stopped-JVP gradient as reference. Oracle scales the same unweighted gradient using the revealed reference. Cosine compares cubic with exact; BA evaluates the sign of projection onto original G_e. n_-/n_+ counts negative/positive truth projections. Each row averages eight measurements from one training seed.

	Relative gradient error ↓				Cosine ↑	BA (%) ↑	n_-/n_+
Seed	Cubic	Unweighted	Scalar	Oracle			
920101	0.4483	1.0101	0.9882	0.9471	0.9185	85.7	1/7
920103	0.4643	1.0036	0.9709	0.8543	0.9253	75.0	4/4
920107	0.4118	1.0193	1.0105	0.9125	0.9000	56.7	3/5

Transport reduces mean relative error by 54.6%, 52.2%, and 59.3%, while projection signs are less accurately recovered (Table 18). All seven sign mismatches occur at JVP-trained checkpoints. Five have exact projection cosines of magnitude at most 0.0411; the other two are -0.0740 and -0.1431, predicted as $+0.0777$ and $+0.0251$. Across the seven cases, cubic–exact gradient cosines remain 0.577–0.942: overall alignment does not ensure the correct component along G_e.

Consequences of reconstructed AdamW steps. Six branches use zero, endpoint, exact stopped-JVP, cubic, unweighted, or same-fit scalar gradients at the stored rate (2×10^{-6} for endpoint-trained models; 5×10^{-7} for JVP-trained models). All gradients remain below the clipping threshold of one. The same 128 held-out samples, separate from updates and shared across seeds, evaluate every branch. Parameters, labels, and loss reduction use FP32; FP32 forwards permit cuDNN TF32. Repeated readouts are identical.

Table 19: **One-step evaluation consequences of reconstructed gradients.** Entries are mean absolute differences in held-out field MSE from the exact stopped-JVP AdamW branch, in units of 10^{-6} (lower is closer). Each mean uses eight batch/checkpoint measurements within a training seed. FP32* allows cuDNN TF32; BF16 uses the forward precision of training. Field errors are measured after executing each update.

Seed	Evaluation forward	Cubic	Unweighted	Scalar
920101	FP32*	0.403	2.307	2.161
	BF16	3.717	4.380	2.751
920103	FP32*	0.804	2.833	2.597
	BF16	4.416	7.446	7.664
920107	FP32*	0.556	1.786	1.837
	BF16	5.153	4.263	5.993

Cubic reconstruction more closely reproduces the exact branch's FP32* evaluation error than the unweighted or scalar fit in all three seeds; the BF16 ordering is mixed (Table 19).

Horizon and state contributions to the learning direction. We separate the returned horizon derivative from state transport on natural errors. Three JVP-trained checkpoints at 3,000 updates

share 16 batches of 64 samples. Path indices are drawn independently with replacement from the last 256 teacher paths, excluded from training. Horizons and times follow the original sampler: $h \sim U[1/8, 1/2]$ and $t \mid h \sim U[0, 1-h]$. Let $G_{e,i}$ and $G_{r,i}$ be the endpoint and stopped-JVP gradients on batch i, before clipping. At the cached endpoint $y = x + hu$, write $v_0 = v_{\mathrm{T}}(t, x)$, $v_1 = v_{\mathrm{T}}(t+h, y)$, $A = \partial_t + v_0 \cdot \nabla_x$, and $e = g - u$. Define

$$\begin{aligned} D &= h\partial_h g - (v_1 - u), \\ S &= -hAg + (v_1 - v_0), \qquad r_{a,b} = e + aD + bS. \end{aligned} \tag{34}$$

The original stopped residual is $r = r_{1,1}$. The two derivatives are evaluated independently. Each detached component uses the same prediction Jacobian and MSE normalization as G_r. For $n = 16$ and $b \in \{e, D, S, r\}$, define the cross-batch work $U_{eb} = [n(n-1)]^{-1} \sum_{i \neq j} \langle G_{e,i}, G_{b,j} \rangle$. Positive work indicates descent along $-G_b$ for the smooth surrogate. D contains all fixed-state horizon orders. With exact teacher endpoints, $D = h\partial_h e$ and $S = -hAe$; the measurements use the original numerical labels.

Endpoint, JVP, and zero-gradient branches use the stored rate 5×10^{-7} and evaluation mode, as during training. Whole-batch gradients and field readouts use original BF16 calculations; weights, labels, and loss reductions use FP32. Independent component pullbacks sum to the native gradient within 0.86% in relative norm. Teacher endpoint velocities use the original FP32 model. The clipping threshold is one. Only the first JVP batch at seeds 920101 and 920103 exceeds it.

Table 20: **Contributions to cross-batch learning.** Means ± paired-jackknife SE. Horizon and state terms oppose endpoint descent. Field columns report actual AdamW changes: positive differences mean higher JVP error than the named comparator.

	Gradient work (10^{-3})				Field-MSE difference (10^{-6})	
Seed	U_{ee}	U_{eD}	U_{eS}	U_{er}	JVP−Endpoint	JVP−Zero
920101	25.41 ± 2.08	-26.17 ± 2.12	-18.85 ± 5.07	-19.68 ± 5.11	3.008 ± 1.370	2.995 ± 1.981
920103	13.24 ± 0.73	-10.96 ± 0.72	-8.85 ± 2.49	-6.57 ± 2.39	-0.070 ± 0.956	0.792 ± 1.336
920107	19.27 ± 1.52	-20.56 ± 1.62	-15.72 ± 4.00	-17.05 ± 4.08	2.429 ± 1.304	1.902 ± 1.625

Paired jackknife deletes one batch's row and column together; SEs describe variability across the sixteen batches from the fixed pool. The three seeds share this bank, with 240 off-diagonal evaluations per seed. Horizon feedback cancels most or all of the positive endpoint work in each model; state transport adds another negative term (Table 20). The adverse mean is shared across different batches. Same-batch work is -11.894 ± 10.028, 0.984 ± 5.485, and -13.734 ± 5.834 (mean ± SE, 10^{-3}), respectively. Actual JVP steps yield higher cross-batch error than endpoint steps in two seeds; their difference is near zero at seed 920103.

C.5 Changing the response of natural errors

We compare changes in horizon response with scalar rate tuning and ordinary multi-horizon supervision, keeping the integrated teacher field fixed.

Reweighting horizon and state terms. The interior gradient is $J_h^\top r_{a,b}$, up to the common MSE factor (Eq. (34)). Relative to stopped JVP at $(1,1)$, $(1,1/2)$ halves S; moving to $(1/2,1/2)$ also halves D, giving the $\lambda = 1/2$ feedback mixture.

For $a > 0$ and $0 \le b \le a$, these residuals retain the same boundary-connected solution. Set $c = b/a$ and $t_0 = t + c(h - h_0)$. Along the corresponding teacher characteristic, $r_{a,b} = 0$ implies

$$e(t, x, h) = \left(\frac{h_0}{h}\right)^{1/a} e\big(t_0, \Phi^{\mathrm{T}}_{t_0,t}(x), h_0\big). \tag{35}$$

Here $t_0 + h_0 \le t + h \le 1$, so the characteristic reaches the specified boundary inside the domain. A zero boundary error gives $e = 0$.

All arms receive the teacher endpoints and velocities required by the asymmetric residual. With cached v_0 and boundary supervision fixed, numerical teacher errors introduce $(a-1)\delta u + (b-a)\delta v_1$ in Eq. (34). We measure the gradient contrast from changing D at two JVP checkpoints and two batches

of eight training states. Refining teacher integration changes this contrast by 2.66×10^{-6}–1.46×10^{-5} of its original norm. The best scalar fit between the $(1, 1/2)$ and $(1/2, 1/2)$ interior gradients leaves 10.5–38.1% relative error. With the shared boundary term included, it leaves 10.4–39.1%.

Table 21: **Changing horizon response trades interior and boundary accuracy.** One paired CIFAR-10 initialization, 200 updates, and 256 held-out states per region. Errors are MSE. Each arm shares the same inputs, boundary loss and teacher quantities. The comparison fixes update count, not compute.

			Field MSE ↓		Differential MSE ↓
Residual	(a, b)	Learning rate	Interior	Boundary	Interior
Stopped JVP	$(1, 1)$	10^{-5}	0.008635	0.012292	0.448979
Reduced state term	$(1, 1/2)$	10^{-5}	0.008056	0.009749	0.387516
Reduced both terms	$(1/2, 1/2)$	10^{-5}	0.008079	0.007036	0.371854
Reduced state term	$(1, 1/2)$	5×10^{-6}	0.010936	0.006560	0.335203
Reduced state term	$(1, 1/2)$	2×10^{-5}	0.005655	0.012411	0.196781
Endpoint	$(0, 0)$	10^{-5}	0.010452	0.002554	0.275228

Training shares the initialization, minibatches and optimizer settings of Table 22. The derivative controls use two stopped JVPs and two ordinary predictions per update; endpoint uses two predictions. All arms also evaluate v_1 once per update on 32 states. The 200 updates take 119–122 seconds for the derivative controls and 61 seconds for endpoint.

At the base rate, reducing S lowers both field errors (Table 21). Halving D at fixed S lowers boundary MSE by 27.8%, while interior MSE rises by 0.29%. Doubling the rate with the larger D gives 30.0% lower interior MSE than reducing both terms, at higher boundary error. Rate tuning gives the best interior accuracy; the tested horizon correction instead favors the boundary.

Horizon averaging and ordinary multi-horizon supervision. At fixed (x, t), define the averaged stopped residual

$$\begin{aligned} r_g(z) &= g(t, x, z) - v_{\mathrm{T}}(t, x) - z\mathcal{D}_s g(t, x, z), \\ (Pr_g)(h) &= \frac{1}{h - h_0}\int_{h_0}^{h} r_g(z)\, dz, \end{aligned} \tag{36}$$

with $h_0 = 1/128$ and continuous extension $(Pr_g)(h_0) = r_g(h_0)$. The teacher field has $r_{u^{\mathrm{T}}} = 0$ at every integration node. For a spatially and temporally constant error $e(z) = z^k a$, the original gain $k + 1$ becomes $\sum_{j=0}^{k}(h_0/h)^j$, approaching one at long horizons. For a general error, the same integral satisfies

$$(Pr_g)(h) = \frac{he(t, x, h) - h_0 e(t, x, h_0) + \int_{h_0}^{h} S(t, x, z)\,\mathrm{d}z}{h - h_0}.$$

Horizon integration leaves endpoint terms and integrated state–time transport. The fixed-endpoint characteristic derivative instead includes both changes (Appendix C.4).

We approximate the integral with three equally weighted, stratified uniform nodes. Only the original prediction at h carries gradients. A control uses ordinary stopped-JVP losses at the same three nodes. Writing J_z for the prediction Jacobian, their gradients, up to the shared MSE factor, are

$$\begin{aligned} G_P = J_h^{\top}\sum_j w_j r_g(z_j), \qquad G_{\mathrm{multi}} &= \sum_j w_j J_{z_j}^{\top} r_g(z_j), \\ G_P - G_{\mathrm{multi}} &= \sum_j w_j (J_h - J_{z_j})^{\top} r_g(z_j). \end{aligned} \tag{37}$$

Both constructions sample shorter horizons; their difference is the Jacobian that returns each residual to parameter space.

All arms start from initialization seed 953101 and share the same 200 minibatches, with 32 interior and 32 boundary states per update. Training uses FP32 raw weights, AdamW, a 100-update warmup, and peak rate 10^{-5}; the half-rate arm uses 5×10^{-6}. The half-residual arm scales only the interior residual, leaving boundary supervision unchanged. The two multi-horizon arms share the independent

Table 22: **Residual averaging and ordinary multi-horizon supervision give similar short-run errors.** One paired CIFAR-10 initialization, 200 updates, and 256 held-out states per region. Each arm uses the same boundary loss. Field and differential columns report MSE; the last column counts ordinary prediction calls and JVP primal calls during training.

	Field MSE ↓		Differential MSE ↓	Prediction/JVP
Supervision	Interior	Boundary	Interior	calls
Endpoint	0.010452	0.002554	0.275228	400/0
Stopped JVP	0.008635	0.012292	0.448979	400/200
Stopped JVP, half learning rate	0.011111	0.007332	0.361405	400/200
Stopped JVP, half residual	0.008265	0.007073	0.383580	400/200
Residual average, original J_h	0.008927	0.004585	0.333944	400/600
Ordinary multi-horizon JVP	0.008965	0.004635	0.333315	800/600

quadrature stream (seed 953701). Endpoint labels serve only the common boundary loss and held-out evaluation; interior targets use teacher velocities at the original states.

Pure-interior gradient cosines between the two constructions are 0.99916–0.99969 at the 200-update JVP checkpoint and 0.99784–0.99833 at an independent 1,000-update JVP checkpoint. Two shared batches of eight states and shared nodes give relative gradient differences of 2.93–4.10% and 7.75–9.20%, without optimizer updates. The different order responses produce similar gradients and short-run field errors (Table 22); ordinary multi-horizon supervision reproduces the changes from residual averaging.

C.6 Matched differential-target controls

Table 23: **Single-pair target comparisons.** Fixed endpoint and stopped-JVP objectives use paired initialization and training states, with separately selected learning rates. Each row evaluates one training pair and one shared generation stream. Rows at the same update count reuse checkpoints across NFE. The iMF-style control adds a diagonal velocity loss and adaptive weighting. CIFAR-10 uses CleanFID; FFHQ-64 uses NVLabs EDM FID. Lower FID is better.

Dataset	Updates	Images	NFE	Endpoint	JVP	iMF-style	JVP−End
CIFAR-10	1,000	2,000	5	**21.338**	32.172	48.711	+10.834
	1,000	2,000	7	**21.264**	37.119	35.011	–
	3,000	50,000	5	**7.572**	17.500	–	+9.928
	3,000	50,000	7	**7.689**	21.681	–	+13.992
FFHQ-64	1,000	2,000	8	**12.444**	42.136	–	+29.692

The 1,000-update CIFAR-10 arms use the same 35.76M-parameter backbone, initialization, minibatch sequence, 1,000-update budget, raw weights, and 2,000-image generation seed. The MeanFlow-JVP row reports the best result over the stated learning-rate grid. The iMF-style row adds a trainable diagonal velocity loss and the published per-sample adaptive weighting to the same backbone.

The FFHQ arms use the same 61.81M-parameter architecture, initialization, cache, minibatch sequence, and evaluation noise. Endpoint training takes 433 seconds and 10.49 GB of peak memory, compared with 3,245 seconds and 24.57 GB for the best stopped-JVP arm.

To test the target gap away from teacher trajectories, we collect the NFE-5 states visited by both 3,000-update students under 512 common noises. From each state, RK4–64 supplies the next teacher endpoint. Table 24 evaluates both predictors on both state distributions; paired intervals use 10,000 bootstrap replicates.

Paired bootstrap intervals for the endpoint–JVP differences exclude zero for all six measurements.

The first-correction law transfers across image domains. We repeat the shrinking-interval measurement with released FFHQ-64 and AFHQv2-64 EDM teachers. Each row uses four paired paths at four noise levels. Fresh RK4 endpoints supply the integrated finite-step target, while a neural JVP supplies the stopped first correction. Figure 3 reports the smallest numerically resolved interval.



Figure 8: **Same-noise FFHQ-64 target comparison.** Both rows use the same initialization, minibatch sequence, 1,000-update budget, Euler-8 solver, and generation seed. Fixed endpoint labels preserve facial structure and natural texture; the best stopped-JVP rate produces visible oversaturation and local texture corruption.

Table 24: Cross-evaluation on student rollout states. Local rows pool all five NFE-5 intervals; composition uses the direct RK4–64 final endpoint.

Evaluation states	Endpoint	JVP	Reduction
Endpoint rollout, local MSE	7.257e−5	7.208e−4	89.93%
JVP rollout, local MSE	1.011e−4	9.921e−4	89.81%
Union, local MSE	8.683e−5	8.565e−4	89.86%
Union, local LPIPS	0.004711	0.014831	68.23%
Final composition MSE	5.833e−3	6.456e−2	90.96%
Final composition LPIPS	0.092366	0.278466	66.83%

Table 25: First-correction geometry under RK4 refinement. Each entry uses the same paths, states, and neural JVP; only the number of teacher microsteps within the short endpoint interval changes.

Dataset / RK4 microsteps	2	4	8	16
FFHQ projection	1.9470	1.9475	1.9464	1.9470
FFHQ cosine	0.9886	0.9888	0.9889	0.9889
FFHQ gap exponent	0.9719	0.9721	0.9723	0.9723
AFHQv2 projection	1.9325	1.9314	1.9311	1.9336
AFHQv2 cosine	0.9856	0.9857	0.9857	0.9857
AFHQv2 gap exponent	0.9674	0.9680	0.9681	0.9675

The projection approaches two on all three datasets, its direction aligns with the endpoint correction, and the residual target gap remains first order in interval length.

Public-scale target stress test. We repeat the target comparison from the same released 56.40M-parameter OpenAI CIFAR-10 initialization. The two arms consume the same cached trajectories, minibatch sequence, initial optimizer state, and NFE-2 generation noises, while each uses a separate learning-rate grid. After 2,000 updates, endpoint regression reaches raw 2,000-image FID 19.724; the best stopped-JVP rate reaches 345.769. Automatic JVPs on the NCSN++ network have cosine 0.9950–0.9995 with centered finite differences, while projected-target LPIPS against the exact cached endpoint grows from 0.014/0.298 at initialization to 0.119/0.776 after training in the two trajectory slots. The larger architecture retains a wide gap between the two targets after this short adaptation.

C.7 Native consistency comparison

We also compare finite-step endpoint and consistency predictors using their native output maps and samplers. Both start from the released 56.40M-parameter OpenAI EDM model and use the same RK4 teacher states, batch size 64, training sequence, LPIPS loss, and RAdam optimizer (Liu et al., 2020). The two intervals follow the noise grid $(80, 0.585348, 0.002)$. Endpoint sampling composes deterministic local maps; consistency sampling uses the published two-call re-noising rule (Song et al., 2023). Learning rates are selected on a 100-update pilot and fixed at 4×10^{-4} for both arms. The 1,000-update runs restart from the released initialization. Evaluation uses raw weights and a fresh common stream of 50,000 initial noises.



Figure 9: **Same-noise public-architecture samples.** Columns compare the released CD-LPIPS model, fixed endpoint regression, and the matched RK4–64 teacher at NFE 1 and 2. Every row begins from the same Gaussian tensors; each method retains its native sampler.

Table 26: **Native two-call generation from a common initialization.** CleanFID on CIFAR-10, with one training seed. Student forward counts are training examples multiplied by calls per example.

Training target	Updates	Examples	Student forwards	FID ↓
Fixed endpoint	1,000	64,000	64,000	6.964
Stopped consistency	1,000	64,000	128,000	**5.960**
Fixed endpoint	2,000	128,000	128,000	6.716

Native consistency retains its advantage at both equal-update and equal-student-forward budgets (Table 26). The comparison changes the predicted object and sampler together with the target. It complements the shared-field endpoint–JVP experiment by showing how a complete consistency method behaves under the same teacher and initialization.

D Additional Proofs

The proofs assume a unique teacher flow with two-parameter composition and square-integrable targets and predictors. The endpoint bound requires Lipschitz teacher step maps on the reachable domain; the expansion requires v_{T} to be C^3 near the relevant characteristics.

D.1 Exact execution and finite-step composition

The definition in Eq. (1) gives the exact map and its squared-error identity:

$$\mathcal{E}_h^{u_h^{\mathrm{T}}}(t,x) = x + h u_h^{\mathrm{T}}(t,x) = \Phi^{\mathrm{T}}_{t\to t+h}(x), \tag{38}$$

$$\mathbb{E}\left\|x + hf(t,x) - \Phi^{\mathrm{T}}_{t\to t+h}(x)\right\|_2^2 = h^2 \left\|f - u_h^{\mathrm{T}}\right\|_{L^2}^2 . \tag{39}$$

Field error becomes displacement error through multiplication by h. If $y = \Phi^{\mathrm{T}}_{t\to t+h}(x)$, splitting the displacement at y and applying the teacher semigroup gives

$$(h+k)u^{\mathrm{T}}_{h+k}(t,x) = h u_h^{\mathrm{T}}(t,x) + k u_k^{\mathrm{T}}\big(t+h, x + h u_h^{\mathrm{T}}(t,x)\big) . \tag{40}$$

Subtracting Eq. (38) gives the exact tangent defect

$$\mathcal{E}_h^{v_T}(t,x) - \Phi_{t\to t+h}^{T}(x) = -h\Delta_h^{T}(t,x). \tag{41}$$

Taylor expansion along a teacher characteristic yields Eq. (8), with the next term $h^2(\partial_t + v_T\cdot\nabla_x)a_T/6$.

D.2 Endpoint error bound

If $\sup_x \left\| f(t_i,x,h_i) - u_{h_i}^{T}(t_i,x)\right\| \le \varepsilon_i$ and the ith teacher step map is L_i-Lipschitz, then trajectories starting from the same state obey

$$\left\|\widehat{X}_N - X_N^{T}\right\| \le \sum_{i=0}^{N-1} |h_i|\varepsilon_i \prod_{j=i+1}^{N-1} L_j. \tag{42}$$

For $\Phi_i^{T} = \Phi_{t_i\to t_{i+1}}^{T}$, Eq. (38) gives

$$\left\|S_i(x) - \Phi_i^{T}(x)\right\| \le |h_i|\varepsilon_i. \tag{43}$$

The Lipschitz condition yields

$$\left\|\widehat{X}_{i+1} - X_{i+1}^{T}\right\| \le |h_i|\varepsilon_i + L_i\left\|\widehat{X}_i - X_i^{T}\right\|. \tag{44}$$

Repeated substitution proves Eq. (42).

D.3 Stopped-target transport and horizon spectrum

First-order parameterization note. For a target fixed point $\mathcal{Q}(g^\star) = g^\star$, the differential of relaxed target fitting is

$$\delta g^+ = \left[I - \alpha\{I - D\mathcal{Q}[g^\star]\}\right]\delta g. \tag{45}$$

An eigenmode of $D\mathcal{Q}[g^\star]$ with eigenvalue μ therefore receives multiplier $1-\alpha(1-\mu)$. External labels have $\mu = 0$; the differential horizon blocks have $\mu_k = -k$.

For a parameterized predictor, freeze the evaluation batch and a parameter-space preconditioner B. Let $(J\delta\theta)(z) = \partial_\theta f_{\theta^\star}(z)\delta\theta$ be the parameter-to-function tangent map, where $z = (t,x,h)$, and set $J^* r = n^{-1}\sum_i \partial_\theta f_{\theta^\star}(z_i)^\top r(z_i)$ on the training batch. The half-squared stopped-loss update is $F(\theta) = \theta - \eta B J_\theta^*[f_\theta - \mathcal{Q}(f_\theta)]$. Its residual vanishes at $\theta^\star$, giving

$$\begin{aligned} DF(\theta^\star) &= I - \eta B J^*\{I - D\mathcal{Q}[g^\star]\}J, \\ \delta f^+ &= \left[I - \eta J B J^*\{I - D\mathcal{Q}[g^\star]\}\right]\delta f. \end{aligned} \tag{46}$$

Here $\delta f = J\delta\theta$ and $\delta f^+ = JDF(\theta^\star)\delta\theta$ are differentials. For finite parameter perturbations, the expansion has a remainder $O(\|\delta\theta\|^2)$ when the maps are twice differentiable. For $\mathcal{Q} = \mathcal{T}_\lambda$, $D\mathcal{Q}[g^\star] = \lambda\mathcal{L}$. The target derivative acts on the smooth perturbation function before its residual is evaluated on the batch; its derivative information is retained. If $K = JBJ^*$ preserves an order-k diagonal mode and acts on it with eigenvalue $\kappa_k > 0$, the local multiplier is $1-\eta\kappa_k(1+\lambda k)$. Kernel mixing couples the other modes. For AdamW, the full update acts jointly on parameters and moment states; its derivative includes moment normalization and weight decay. Appendix C.3 specifies the actual-state AdamW measurements in Figure 5.

Along the characteristic with fixed $s = t+h$, $\Phi_{t\to s}^{T}(x) = x + hu_h^{T}(t,x)$ is constant. Differentiating gives $v_T - u_h^{T} + h\mathcal{D}_s u_h^{T} = 0$, so $u_h^{T} = \mathcal{T}u_h^{T}$. Subtracting this equality from $\mathcal{T}g$ proves Eq. (6).

For the jet theorem, work on $\mathcal{J}_K(V) = V^{K+1}$ with $\mathcal{A}(h) = \sum_{j\ge 0} h^j A_j$ and $A_j : V \to V$. Then $h\mathcal{A}(h)$ raises horizon degree whereas $-h\partial_h(h^k e_k) = -kh^k e_k$. The truncated operator is therefore block triangular with diagonal blocks $0, -I, -2I, \ldots, -KI$. Restricting to $e(0) = 0$ removes the first block. This proves Eq. (10) and its boundary-anchored form. A relaxed update has diagonal factors $1-\alpha(k+1)$; replacing the self-dependent part by weight λ changes them to $1-\alpha(1+\lambda k)$ and gives Eq. (11). The degree-raising blocks can produce transient amplification, but a finite triangular update whose diagonal multipliers lie strictly inside the unit disk is asymptotically stable.