LET IT BE SIMPLE: ONE-STEP ACTION GENERATION FOR VISION-LANGUAGE-ACTION MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

One-step generation from standard diffusion or Flow-Matching training is rarely effective for text-to-image or video generation without distillation or specialized one-step objectives. We find that VLA action generation behaves strikingly differently. Biasing only the training timestep distribution toward the noise endpoint yields strong one-step policies across LIBERO, LIBERO-Plus, LIBERO-Pro, and real-world robot tasks. Moreover, for open-source VLA models with different architectures, simply replacing iterative decoding with a single forward pass often performs comparably to the original multi-step decoding on benchmarks such as LIBERO, CALVIN, and SimplerEnv.

We attribute this difference to the condition-target structure of VLA generation, which is closer to image-to-text than text-to-image: rich observations strongly constrain a comparatively compact target. We formalize this intuition through the irreducible velocity loss $R_v(t, c)$ of standard Flow Matching, which becomes much smaller toward the noise endpoint for VLA, whereas image-generation models usually exhibit a U-shaped profile.

The geometry induced by the learned generative map shows the same contrast: the Jacobian of this map with respect to the initial noise has a substantially smaller largest singular value for action generation than for image generation. Together, these results suggest that VLA's condition-target structure makes one-step generation unusually accessible to standard Flow Matching.

1 INTRODUCTION

Diffusion and flow models typically trade inference latency for sample quality via iterative numerical integration (Ho et al., 2020; Albergo et al., 2025; Lipman et al., 2023). In text-to-image and video generation, few-step or single-step sampling remains notoriously challenging because a sparse text prompt leaves a broad, high-dimensional, and multimodal target distribution (Song et al., 2023; Pham et al., 2026). Consequently, effective fast generators in vision have historically necessitated specialized machinery, such as consistency training, adversarial distillation, flow-map, or mean-flow formulations (Song et al., 2023; Yin et al., 2024; Frans et al., 2025; Geng et al., 2025; Boffi et al., 2025; Lu et al., 2026). In contrast, we discover that continuous Vision-Language-Action (VLA) flow policies (Black et al., 2025b;a; Luo et al., 2026; Zheng et al., 2025) trained under vanilla conditional flow matching natively achieve remarkable one-step performance. Across diverse architectures and simulation and hardware benchmarks (Liu et al., 2023a; Fei et al., 2025; Zhou et al., 2026; Li et al., 2024; Mees et al., 2022), a single forward pass matches or even outperforms multi-step ODE decoding, with high-noise training schedules further amplifying this capability without altering the objective.

We attribute this fundamental divergence to the underlying *condition-target structure*, which governs the complexity of the residual generative problem after conditioning. In VLA action generation, the policy receives rich multimodal streams—multi-view visual observations, natural language instructions, and proprioceptive state vectors (Black et al., 2025b; Chi et al., 2025; Octo Model Team et al., 2024)—to generate a short action trajectory. From this perspective, VLA decoding aligns with image-to-text transcription rather than open-ended visual synthesis: extensive conditioning tightly constrains the admissible trajectory manifold to a concentrated conditional posterior. Diagnostic ablations on LIBERO confirm this distinction: expanding the action horizon increases target com-

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.

plexity, which multi-step solvers readily decode; and weakening the conditioning signals disrupts the underlying condition structure, degrading closed-loop performance in ways that multi-step sampling cannot rescue.

To formalize this intuition, we analyze the irreducible velocity risk of conditional flow matching (Lipman et al., 2023; Albergo et al., 2025; Liu et al., 2023b). Across diverse architectures and benchmarks, VLA policies consistently exhibit a monotonic loss profile that drops toward the noise boundary, unlike the U-shaped barrier in visual synthesis (Li & He, 2026; Lu et al., 2026). A local geometric lens corroborates this contrast: while visual flows are expansive, VLA transport is predominantly contractive. Yet this contraction does not collapse the policy into deterministic regression: the learned flow retains sparse, noise-dependent spectral fractures—a structure that pure-noise endpoint training fails to recover.

We summarize our primary contributions as follows:

- **Cross-Benchmark Single-Step Parity**: Across standard LIBERO (Liu et al., 2023a), LIBERO-Plus (Fei et al., 2025), LIBERO-Pro (Zhou et al., 2026), CALVIN (Mees et al., 2022), SimplerEnv (Li et al., 2024), and real-world DualYam experiments (Zhang et al., 2026), standard flow-matching policies match multi-step numerical integration in a single step without distillation. Biasing the training timestep schedule toward high-noise states further eliminates residual discretization gaps without altering the objective.
- **Irreducible Velocity Risk and Gaussian Transport Concentration**: We formalize the condition-target dichotomy via the irreducible velocity risk $R_v(t, c)$, proving the endpoint identity $R_v(0, c) = \operatorname{Tr} \operatorname{Cov}(x_1 \mid c)$. Through the pullback metric of the generative map, we establish via the Gaussian Poincaré inequality and concentration bounds that transport contractivity enforces exponential sub-Gaussian concentration of the target around its conditional mean. Empirically, real VLA policies exhibit a monotonic risk profile that vanishes toward the noise endpoint, contrasting sharply with the severe U-shaped barrier of visual synthesis.
- **Transport Geometry and Physical Phase Boundaries**: Evaluating the noise-to-action Jacobian $J = \nabla_{x_0} \phi_1^\theta$ demonstrates that while image flows operate in an expansive regime, VLA policies exhibit a predominantly contractive spectrum ($\sigma_1(J) \leq 1$). Crucially, the sparse spectral fractures where contractivity breaks down demarcate discrete physical phase boundaries, causally steering action switches and reliably predicting gripper actuation timing.

2 Related Work

Fast sampling in flow and diffusion: from vision to action policies. Sampling in diffusion and continuous flow models natively relies on iterative numerical integration (Ho et al., 2020; Song et al., 2021; Lipman et al., 2023; Liu et al., 2023b). In visual synthesis, single-step sampling is notoriously challenging due to the under-constrained multimodal targets, prompting specialized acceleration methods such as distillation (Salimans & Ho, 2022; Yin et al., 2024), consistency models (Song et al., 2023), flow maps (Boffi et al., 2025; Frans et al., 2025), and mean-flow or drift formulations (Geng et al., 2025; Deng et al., 2026). While recent robotics efforts have similarly turned to consistency training or self-distillation to accelerate continuous VLA policies (Song et al., 2025; Luan et al., 2026), we find that standard flow matching natively achieves single-step parity across architectures and benchmarks—a capability further boosted by high-noise training schedules (Esser et al., 2024; Li & He, 2026) without auxiliary machinery.

Condition-target structure across visual and robotic generation. The tractability of conditional generation is governed not merely by raw information quantity, but fundamentally by the *representational structure* of the condition relative to the target manifold. In visual synthesis, conditioning has progressively evolved from discrete class tokens in DiT and SiT (Peebles & Xie, 2023; Ma et al., 2024) and text embeddings (CLIP, T5) (Radford et al., 2021; Raffel et al., 2020) to structured multimodal representations (Kong et al., 2024; Bai et al., 2025; Chen et al., 2026), alongside compact target latent spaces (Rombach et al., 2022; Dieleman, 2025; Yao et al., 2025; Tian et al., 2024; Yu et al., 2023; 2024). Unlike visual synthesis, VLA policies (Black et al., 2025b;a; Luo et al., 2026; Zheng et al., 2025; Chi et al., 2025; Octo Model Team et al., 2024) leverage rich multimodal context to

predict an intrinsically compact action trajectory. Consequently, grounded by structured multimodal conditions, this compact action manifold allows standard flow matching to decode effectively in a single step.

Irreducible velocity risk and Jacobian spectral diagnostics. Time-dependent vector field optimization in flow matching is fundamentally governed by the Bayes conditional variance along the probability path (Lipman et al., 2023; Albergo et al., 2025; Liu et al., 2023b). In diffusion and flow matching, an irreducible velocity risk barrier emerges when the conditioning signal fails to resolve uncertainty over an expansive target space, as observed in under-constrained pixel diffusion (Li & He, 2026; Lu et al., 2026). In contrast, when the condition structure is rich relative to the target—or when the target manifold is intrinsically compact, as in VLA policies—this ambiguity collapses, and the resulting velocity loss is empirically monotonic. Furthermore, while sensitivity and spectrum analysis of generative maps $J = \nabla_{x_0} \phi_1^\theta$ have been studied for controllable editing (Ben-Hamu et al., 2024; Chen et al., 2024) and out-of-distribution detection (Shoushtari et al., 2026), we extend this geometric lens to continuous action policies. Remarkably, this unifying perspective reveals that just as Jacobian singular values capture discrete semantic transitions in visual generation, their spectral ridges in VLA directly demarcate physical phase boundaries—such as contact transitions and gripper actuations—while confirming the contractive nature of action flows.

3 From Velocity Risk to Gaussian Transport Concentration

3.1 Flow Matching Preliminaries

Let $\mathbb{R}^d$ denote the data space. Flow matching (Lipman et al., 2023; Liu et al., 2023b; Albergo et al., 2025) models continuous generative trajectories by regressing a time-dependent vector field $v_\theta(x, t, c)$ to generate a probability path p_t bridging a prior $p_0 = \mathcal{N}(0, I)$ at $t = 0$ to the data distribution p_1 at $t = 1$ via the ODE $\frac{d}{dt}\phi_t(x) = v_t(\phi_t(x))$. Because regressing the marginal vector field $\mathbb{E}_{t,p_t}[\|v_\theta - u_t\|^2]$ is generally intractable, Conditional Flow Matching (CFM) trains the model on conditional probability paths $p_t(x \mid x_1)$ generated by per-sample conditional vector fields $u_t(x \mid x_1)$, which yields identical gradients with respect to θ.

In practice, we adopt the standard linear interpolation $x_t = tx_1 + (1-t)x_0$ with $x_0 \sim \mathcal{N}(0, I)$, which yields a constant conditional velocity $u_t(x_t \mid x_1) = x_1 - x_0$. The policy is trained by optimizing $\mathcal{L}_{\text{CFM}}(\theta) = \mathbb{E}_{t\sim\mathcal{U}[0,1],\, x_0\sim\mathcal{N}(0,I),\, (c,x_1)\sim q} \left\|v_\theta(x_t, t, c) - (x_1 - x_0)\right\|^2$.

3.2 Irreducible Velocity Risk

Under the linear interpolation path, the flow matching loss is lower-bounded by the Bayes conditional variance: $R_v(t, c) \triangleq \mathbb{E}\,\mathrm{Tr}\,\mathrm{Var}(x_1 - x_0 \mid x_t, c) = \frac{1}{(1-t)^2}\mathbb{E}\,\mathrm{Tr}\,\mathrm{Cov}(x_1 \mid x_t, c)$. At $t = 0$, this simplifies to the target dispersion $R_v(0, c) = \mathrm{Tr}\,\mathrm{Cov}(x_1 \mid c)$. Since the endpoint field targets the conditional mean, $v^*(x_0, 0, c) = \mathbb{E}[x_1 \mid c] - x_0$, a single Euler step predicts $\mathbb{E}[x_1 \mid c]$ with expected distortion $R_v(0, c)$. Single-step viability thus hinges on whether conditioning concentrates the posterior measure sufficiently to render this mean a faithful trajectory proxy.

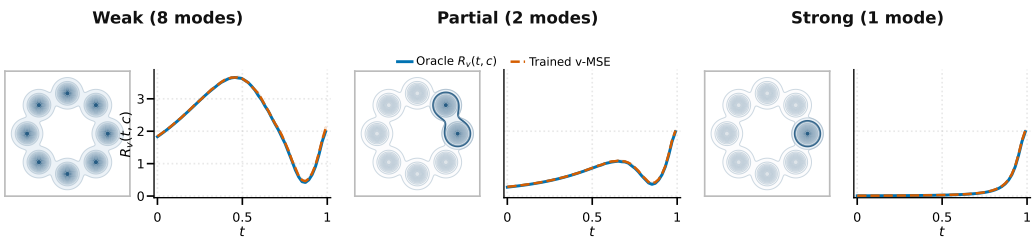

Figure 1: Each column shows the target geometry (left) and the analytical/trained risk profiles (right), from weak to strong conditioning.

For targets distributed as an M-component Gaussian mixture $x_1 \sim \sum_k \pi_k \mathcal{N}(\mu_k, \sigma^2 I)$, the law of total variance yields a closed-form posterior covariance along the linear-Gaussian path:

$$\mathrm{Cov}(x_1 \mid x_t, c) = \frac{\sigma^2(1-t)^2}{(1-t)^2 + t^2\sigma^2} I + \sum_{k \in \mathcal{S}_c} w_k(x_t)(\mu_{k,t} - \bar{\mu}_t)(\mu_{k,t} - \bar{\mu}_t)^\top,$$

where $\mathcal{S}_c$ denotes the subset of modes consistent with condition c, $w_k(x_t) = p(k \mid x_t, c)$ is the posterior component probability, and $\mu_{k,t} = \mathbb{E}[x_1 \mid x_t, k]$.

Figure 1 visualizes this contraction: across all three regimes, an identical MLP tracks the analytical Bayes-risk profile, showing that stronger conditioning directly lowers endpoint velocity risk.

3.3 Pullback Geometry and Gaussian Concentration

While the scalar trace $R_v(t, c)$ tracks aggregate posterior variance, it inherently overlooks the anisotropic geometry of the conditioned target. To resolve this, we leverage the differential properties of the transport map $\phi_1(\cdot; c)$ as an intrinsic probe of target manifold regularity.

The Euclidean metric on the target action space induces a Riemannian pullback metric $g(x_0; c) \triangleq J(x_0; c)^\top J(x_0; c)$ on the prior noise domain via the generative map ϕ_1, where $J(x_0; c) \triangleq \nabla_{x_0} \phi_1(x_0; c)$. The operator norm $\sigma_1(J) = \sqrt{\lambda_{\max}(g)}$ quantifies the maximal local expansion under push-forward. Applying the Gaussian Poincaré inequality $\mathrm{Var}_\gamma[f] \le \mathbb{E}_\gamma[\|\nabla f\|^2]$ coordinate-wise under the prior $\gamma = \mathcal{N}(0, I_d)$ directly bounds the global target dispersion via

$$\mathrm{TrCov}(x_1 \mid c) \le \sum_{i=1}^{d} \mathbb{E}_{x_0}[\|\nabla_{x_0} \phi_1^{(i)}(x_0; c)\|^2] = \mathbb{E}_{x_0}[\|J(x_0; c)\|_F^2] \le d\, \mathbb{E}_{x_0}[\sigma_1(J(x_0; c))^2],$$

so an on-average contractive map ($\mathbb{E}[\sigma_1^2] \le 1$) directly guarantees a low-variance target by compressing the aggregate trace covariance relative to the prior.

While this trace bound controls average dispersion, a uniform spectral bound $L(c) \triangleq \operatorname{ess\,sup}_z \|J(z; c)\|_2$ further upgrades this control to exponential concentration. Via Gaussian concentration (Herbst argument; Appendix A), we obtain

$$\mathbb{P}\left(\|x_1 - \mathbb{E}[x_1 \mid c]\|_2 \ge \sqrt{\mathbb{E}_z \|J(z; c)\|_F^2} + \epsilon \,\middle|\, c \right) \le \exp\left(-\frac{\epsilon^2}{2L(c)^2} \right), \quad \forall \epsilon > 0. \tag{1}$$

Here $\sqrt{\mathbb{E}_z \|J(z; c)\|_F^2} \le L(c)\sqrt{d}$ controls the characteristic radius, while $L(c)$ controls the exponential tail; when $L(c) \le 1$, all centered one-dimensional projections are sub-Gaussian with variance proxy at most 1.

4 Continuous Vision-Language-Action Flow Policies

4.1 Controlled Limit Case: Deterministic Target with Complex Condition

To verify that low-variance target manifolds natively admit single-step flow matching before confronting robotics dataset, we first probe a synthetic limit case: **the target posterior collapses to a Dirac delta, while the condition requires parsing high-dimensional visual scenes.** We instantiate this by mapping an image of a multi-digit MNIST grid (c) to its deterministic token sequence (x_1) using a standard encoder–decoder flow model (Fig. 2(a); details in Appendix E).

Consistent with our risk analysis ($R_v(0, c) \to 0$), empirical velocity MSE minimizes at the noise boundary $t = 0$ (Fig. 2b). Biasing training time toward noise ($t \sim \mathrm{Beta}(0.5, 1.0)$) lifts one-step decoding to 10-step Euler parity—indeed, pure-noise training succeeds on MNIST (Fig. 2c) but fails entirely on CIFAR-10 (Fig. 2d), confirming that target posterior concentration natively enables one-step flow matching.

This limit case also dictates our architecture. Encoder–decoder scaling ablations show that a weaker vision encoder degrades one-step performance, while expanding decoder capacity fails to compensate (Appendix Fig. 14). Thus, the bottleneck lies strictly in condition representation—motivating a powerful vision-language backbone paired with a lightweight flow action head.

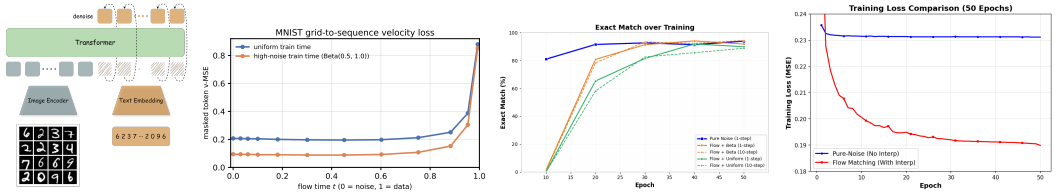

(a) Task & model. (b) v-loss profile. (c) One-step exact match. (d) CIFAR-10 contrast.

Figure 2: Diagnostics on image-to-text MNIST (a–c) and class-to-image CIFAR-10 (d). (a) Input grid. (b) Empirical velocity MSE over flow time t. (c) Exact-match accuracy across training time distributions (uniform, high-noise Beta, and pure noise). (d) Pure-noise training loss contrast.

4.2 Embodied Policy Architecture and Action Flow Parameterization

Robot manipulation operates in this same favorable regime: rich multimodal context (camera streams, language commands, and proprioceptive state vectors) tightly bounds an intrinsically compact target—a short action chunk $x_1 \in \mathbb{R}^{H \times d_a}$ over horizon H.

Architecture Guided by Condition Dominance. Grounded by our finding that generative performance is bottlenecked by condition representation rather than target-side capacity, our architecture decouples high-capacity multimodal perception from lightweight action decoding (Fig. 3). Built upon OpenPI (Black et al., 2025b), a vision-language backbone processes multi-camera observations via SigLIP (Zhai et al., 2023) and fuses them with instruction tokens through PaliGemma (Beyer et al., 2024) to produce contextual prefix tokens. A lightweight bidirectional Transformer head (Luo et al., 2026) takes these prefix tokens along with the robot proprioceptive state s_t, flow time t, and noised action candidate $x_t \in \mathbb{R}^{H \times d_a}$ to predict velocity $v_\theta(x_t, t, c)$.

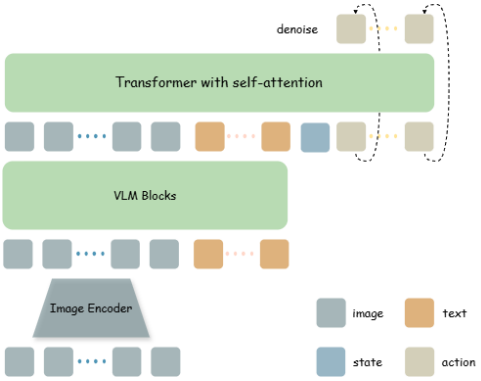

Figure 3: Overview of the flow-based VLA architecture: multimodal condition tokens modulate a lightweight action head that regresses the action trajectory velocity field.

We evaluate across two model capacities while fixing the action head at ∼32M parameters: a *Tiny* variant (∼109M active parameters) for extensive diagnostic sweeps, and a *Full* variant (∼888M active parameters) to examine scaling behavior (architectural breakdown in Appendix F).

5 Robot Manipulation Experiments

5.1 High-Noise Training Unlocks One-Step VLA Policies

Evaluating noise-shifted training ($t = u/[1+(\alpha-1)(1-u)], \alpha \geq 1$) across LIBERO suites (Liu et al., 2023a) (Table 1) reveals that **biasing training toward noise** ($\alpha = 4$) **eliminates the one-step gap**, consistently outperforming the 10-step uniform baseline across all four suites. Naturally, over-weighting the noise endpoint sacrifices intermediate path accuracy, degrading 10-step integration under extreme bias ($\alpha = 8$). Indeed, whereas image flow matching hinges on noise-data interpolation to optimize complex distributions, the success of

Table 1: **Time-schedule controls on LIBERO** (success rates %). Differences (±) are relative to Uniform 1-step; best 1-step results are **bolded**.

Schedule	Steps	Spatial	Object	Goal	Long
Uniform	10	96.6	96.2	93.2	80.8
Uniform	1	88.8	92.8	90.2	70.2
$\alpha = 3$	1	95.8 +7.0	97.0 +4.2	94.4 +4.2	78.0 +7.8
$\alpha = 3$	10	91.6	96.2	82.2	52.4
$\alpha = 4$	1	**96.4** +7.6	**99.6** +6.8	**96.8** +6.6	**85.2** +15.0
$\alpha = 4$	10	93.4	96.8	91.4	63.4
$\alpha = 8$	1	95.6 +6.8	99.0 +6.2	94.0 +3.8	78.0 +7.8
$\alpha = 8$	10	78.4	40.0	56.6	22.6
Pure-noise	1	96.0 +7.2	97.8 +5.0	96.2 +6.0	72.8 +2.6

pure-noise training highlights that strong conditioning places embodied control in a fundamentally different geometric regime.

Encoder Scaling and the Necessity of Prior Geometry. Scaling the perceptual encoder to the Full variant pushes one-step decoding to full parity with 10-step Euler integration across all suites (Table 2). Yet, we also try the direct regression: training a Tiny variant that completely replaces prior noise with learnable input embeddings degrades substantially on complex suites (Appendix Table 6). This gap indicates that the target distribution retains nontrivial structure, motivating our geometric analysis in Section 3.3.

Table 2: **Full-encoder scaling performance** ($\alpha = 3$, success rates %). One-step decoding matches 10-step Euler integration across all suites. Differences ($\pm$) are relative to 10-step.

Steps	Spatial	Object	Goal	Long
1-step	98.4 (-0.4)	100.0 **+1.0**	97.0 (-1.0)	95.6 **+0.6**
10-step	98.8	99.0	98.0	95.0

Generalization Across Diverse Policies. This behavior also holds across diverse public checkpoints(Zheng et al., 2025; Ye et al., 2026; Black et al., 2025b;a), where one-step decoding consistently matches multi-step integration across LIBERO (Table 3), SimplerEnv WidowX (Li et al., 2024), and CALVIN (Mees et al., 2022) (Table 4). This suggests that near-lossless single-step execution is an inherent property of strongly conditioned action flows. Single-step decoding also cuts per-query model latency 1.5–3.4× across architectures, with the condition prefix dominating the remainder (Appendix G.1).

Table 3: **Paired decoding of official π checkpoints on LIBERO** (success rates %, 500 episodes/suite). Differences ($\pm$) are relative to each checkpoint's 10-step baseline.

Model	Steps	Spatial	Object	Goal	Long
$\pi_{0.5}$	10	99.0	99.4	98.0	94.2
$\pi_{0.5}$	1	99.4 **+0.4**	99.4 (0.0)	98.8 **+0.8**	95.6 **+1.4**
π_0	10	95.4	97.6	96.0	85.6
π_0	1	95.4 (0.0)	97.6 (0.0)	96.0 (0.0)	82.2 **-3.4**

Table 4: **Cross-benchmark paired decoding of third-party checkpoints.** Evaluation on SimplerEnv WidowX (Li et al., 2024) (24 episodes/task) and CALVIN ABC→D (Mees et al., 2022) (subgoal success %; Avg. chain length / 5). Differences ($\pm$) are relative to each checkpoint's multi-step baseline.

	SimplerEnv WidowX					CALVIN ABC→D					
Model	**Spoon**	**Carrot**	**Cube**	**Eggpl.**	**Avg.**	**LH-1**	**LH-2**	**LH-3**	**LH-4**	**LH-5**	**Avg.**
Qwen2.5-π (4-step)	79.2	58.3	37.5	83.3	**64.6**	91.4	83.0	74.6	67.0	59.3	**3.75**
Qwen2.5-π (1-step)	91.7 **+12.5**	58.3 (0.0)	33.3 **-4.2**	70.8 **-12.5**	**63.5** (-1.1)	91.9 **+0.5**	83.7 **+0.7**	74.7 **+0.1**	66.5 (-0.5)	57.5 **-1.8**	**3.74** (-0.01)
X-VLA (10-step)	95.8	91.7	87.5	87.5	**90.6**	95.6	91.8	87.1	81.7	73.9	**4.30**
X-VLA (1-step)	100.0 **+4.2**	100.0 **+8.3**	87.5 (0.0)	87.5 (0.0)	**93.8 +3.2**	95.0 (-0.6)	91.5 (-0.3)	86.4 (-0.7)	80.6 (-1.1)	72.4 **-1.5**	**4.26** (-0.04)

Real-World Bimanual Manipulation. We further validate one-step decoding on physical hardware using the dual-arm Yam platform across three tasks (Zhang et al., 2026) (Fig. 15). For each task, we fine-tune a policy from OpenPI $\pi_{0.5}$ (Black et al., 2025b) using 10 hours of human teleoperation demonstrations. One-step decoding matches or rivals the official 10-step baseline (Table 5), demonstrating real-world viability.

Table 5: **Real-world bimanual task success** (success / trials, 15 trials/task; Fig. 15). 10-step results are from the workshop official baseline.

Task	1-step	10-step
Insert mouse battery	80.0%	80.0%
Seal bottle cap	46.7%	35.0%
Tower of Hanoi	40.0%	50.0%

We show this parity persists under perturbations: one-step decoding consistently matches or even outperforms multi-step integration across LIBERO-Plus(Fei et al., 2025) and LIBERO-Pro(Zhou et al., 2026) (Appendix H; numerical solver audit in Appendix Fig. 16).

5.2 Velocity Loss Landscape Across Flow Time

Across diverse VLA architectures and benchmarks (Fig. 4a–g), velocity MSE increases monotonically from the noise to the clean target, matching our $R_v(t, c)$ formulation (Section 3). The CIFAR-10 ablation (Fig. 4h) explains this transition: while weak class conditioning produces a U-shaped profile, progressively adding low-resolution spatial context flattens the curve into the same monotonic profile observed in VLAs. Furthermore, scaling from the Tiny to the Full encoder uniformly suppresses the

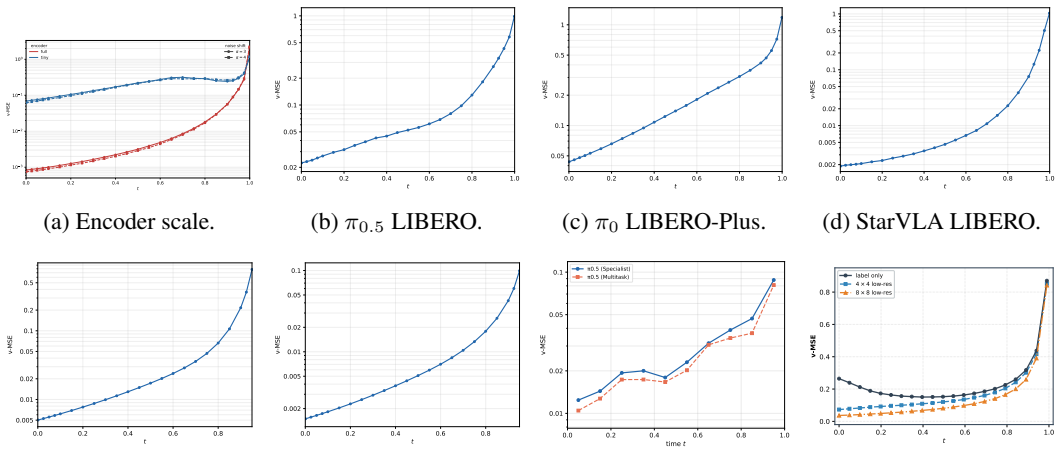

(a) Encoder scale. (b) $\pi_{0.5}$ LIBERO. (c) π_0 LIBERO-Plus. (d) StarVLA LIBERO.

(e) X-VLA CALVIN. (f) X-VLA WidowX. (g) $\pi_{0.5}$ real robot. (h) CIFAR-10.

Figure 4: **Fixed-time velocity MSE across flow time** $t \in [0, 1]$. (a–g) Robotic policies across diverse backbones and benchmarks. (h) CIFAR-10 under varying conditioning (class label vs. 4×4 / 8×8 low-resolution context). Scaling the vision-language backbone uniformly lowers the error profile ((a), Tiny vs. Full).

entire loss curve (Fig. 4a), confirming that capacity scaling directly reduces vector field residual to boost single-step performance.

5.3 Expanding Action Horizons and Single-Step Viability

We test the limits of single-step viability by extending the action horizon from $H = 10$ to $H \in \{20, 30, 40\}$ on LIBERO-Long (Fig. 5b). As the chunk length grows, single-step decoding degrades systematically: while single-step execution dominates at standard horizons ($H = 10$), multi-step integration overtakes it at larger H, with noise shifting ($\alpha = 4$) delaying this crossover.

While velocity profiles preserve the familiar early-time minimum across all H (Fig. 5a), loss magnitudes do not cleanly rank with horizon length. Although velocity loss serves as a useful proxy for endpoint risk, it conflates optimization error with action dimensionality across horizons. We complement it with the local Jacobian spectrum (σ_1) to examine the flow sensitivity directly.

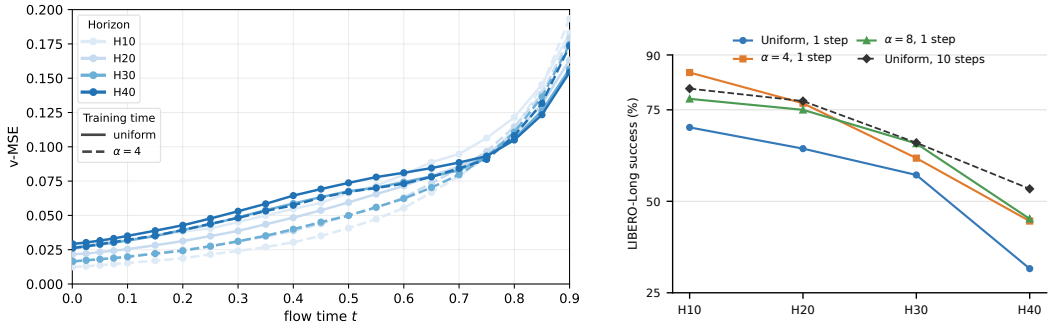

Figure 5: **Action horizon scaling on LIBERO-Long.** (a) Velocity MSE profiles under uniform and noise-shifted schedules across horizons $H \in \{10, 20, 30, 40\}$. (b) Single-step advantage narrows as horizon grows, with multi-step integration overtaking it at larger H.

6 Empirical Transport Geometry: Global Contraction and Local Phase Boundaries

6.1 Global Contraction of the Transport Map

While our bounds link single-step viability to posterior compactness, evaluating target dispersion directly from offline demonstrations is impossible, as each context $c = (o_t, l, s_t)$ typically admits only a single trajectory. A trained flow model bridges this gap: by providing a smooth generative map $x_1 = \phi_1^\theta(x_0; c)$ from prior noise to actions, its Jacobian $J(x_0; c) \triangleq \nabla_{x_0}\phi_1^\theta(x_0; c)$ acts as an exact differential probe of how conditioning contracts the target measure.

To capture this behavior across the entire action manifold, we evaluate the empirical population distribution of the spectral norm $\sigma_1(J(x_0; c)) = \|\nabla_{x_0}\phi_1^\theta\|_2$ by jointly sampling condition contexts from evaluation demonstrations and prior noise $x_0 \sim \mathcal{N}(0, I)$. Crucially, rather than characterizing an isolated context, this empirical distribution serves as a dataset-level diagnostic of typical transport geometry across tasks. To probe the true transport geometry, we evaluate $J(x_0; c) = \nabla_{x_0}\phi_1^\theta$ where $\phi_1^\theta(x_0; c) \triangleq x_0 + \int_0^1 v_\theta(x_t, t, c)\, dt$ for flow-matching model, and $\phi_1^\theta(x_0; c) \triangleq G_\theta(x_0, c)$ for native one-step models (MeanFlow (Geng et al., 2025), pMF (Lu et al., 2026)).

As shown in Figure 6, the cross-domain divide is unequivocal: singular values for VLA policies concentrate predominantly below the contractive threshold $\sigma_1 = 1$, a predominantly contractive regime consistent with the concentration behavior of Eq. (1), whose exponential guarantee itself rests on the stronger uniform bound $L(c)$. In contrast, image flows are strictly expansive, operating two to three orders of magnitude higher. Within the VLA family, extending the action horizon progressively shifts mass past $\sigma_1 = 1$, capturing the loss of contractivity that underpins the performance degradation in Section 5.3.

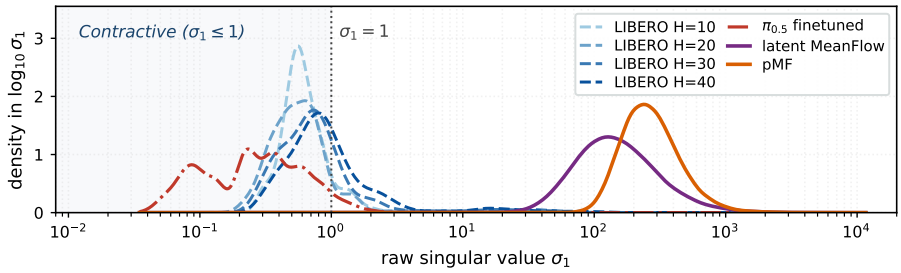

Figure 6: **Empirical distribution of the generative Jacobian spectral norm $\sigma_1(J)$.** KDE densities over $\log_{10}\sigma_1$ across domains. Output conventions cannot account for this gap, as variance-matching the $[-1, 1]$ image coordinates would only enlarge the image spectra (Appendix C.3)

6.2 Local Geometry: σ_1 Fractures Trace Decision Boundaries

Beyond aggregate contractivity (Section 6.1), $\sigma_1(x_0; c)$ measures local sensitivity to initial noise perturbations. Mapping a smooth Gaussian prior onto separated data modes forces the transport map to stretch at the interfaces, which manifest as high-σ_1 ridges separating outcome basins. Throughout this section we write the Jacobian SVD as $J = \sum_i \sigma_i\, u_i v_i^\top$, with v_i acting in noise space and u_i in action space. We probe this structure identically across domains: fix the condition, sample Gaussian noise vectors, screen for high-σ_1 anchors, and map σ_1 on the plane spanned by each anchor's leading right singular directions (v_1, v_2), decoding along a circular orbit that crosses the ridge.

We first inspect visual synthesis, where mode boundaries are semantic transitions. Fixing a class label of the MeanFlow model, a prominent singular-value ridge emerges in the (v_1, v_2) plane (Figure 7a); crossing it on the orbit abruptly flips the mantis pose between samples #6 and #7, while poses on either side remain stable (Figure 7b). The flow thus maps mode adjacency back into noise space, where σ_1 exposes the interfaces between coherent semantic modes.

In flow policies, we sample context–noise pairs jointly: contexts $c = (o_t, l, s_t)$ from closed-loop episodes, $x_0 \sim \mathcal{N}(0, I)$, screened by $\sigma_1(x_0; c)$. The high-σ_1 pairs cluster near *gripper transitions*—steps at which the decoded chunk switches the gripper—while σ_1 stays below 1 almost everywhere

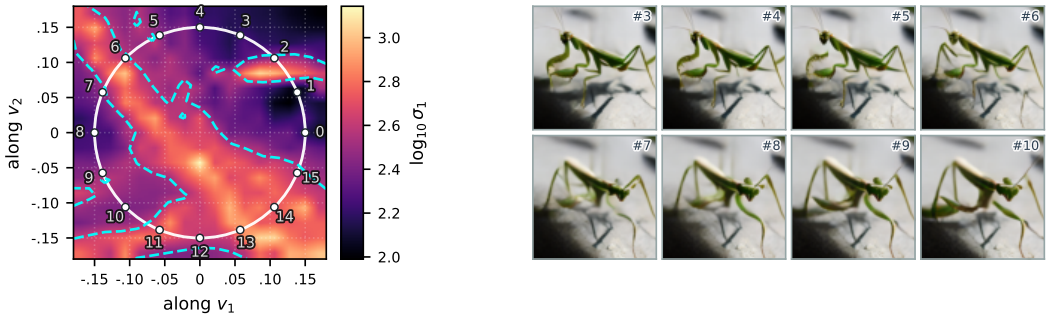

Figure 7: **Local transport geometry of an image flow.** (a) $\log_{10}\sigma_1$ over the (v_1, v_2) noise plane; the circular orbit (cyan) crosses the p95 ridge. (b) Generations along the orbit; crossing the ridge flips the pose between #6 and #7.

else (Fig. 8a), keeping the rollout concentrated around the conditional mean (Eq. (1); protocols in Appendix C).

Write $m(x_0; c) \in \{1, \ldots, H\}$ for the first chunk step of $\phi_1^\theta(x_0; c)$ in the new gripper state. At an anchor near such a transition, a sharp σ_1 ridge emerges in the (v_1, v_2) plane, and crossing it shifts m exactly once (Fig. 8b,c).

Two checks establish this ridge as a physical decision boundary: (i) **Causal specificity:** Pushing initial noise along v_1 ($\epsilon \leq 0.1$) reliably shifts the switch step of screened events in the top σ_1 decile (by at most two chunk steps), whereas a pure-noise policy responds in only 6.5% of cases (Appendix C, Fig. 10a). (ii) **Predictive timing:** The gripper-channel singular peak $h^* \triangleq \mathrm{argmax}_h(u_1^{\mathrm{grip}}[h])^2$ falls within two chunk steps of the *recorded* gripper transition at 81% of anchors (median offset -1; Fig. 8d).

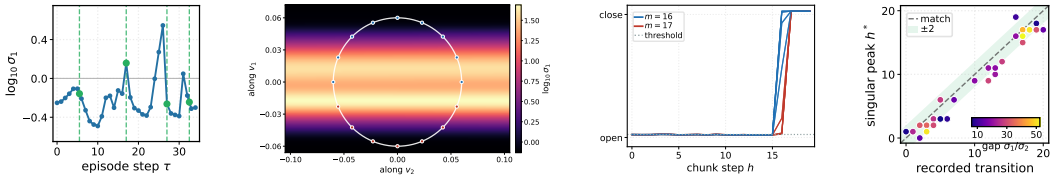

(a) episode σ_1 trace (b) $\log_{10}\sigma_1$ on (v_1, v_2) (c) orbit decodes (d) h^* vs. recorded

Figure 8: **Action flow geometry during manipulation.** Protocols and additional panels: Appendix C. Together, these views reconcile: global contractivity concentrates one-step samples near the conditional mean via Gaussian concentration (Eq. (1)), while spectral fractures isolate discrete physical decisions. Even under single-step decoding, this boundary timing signature is preserved far more accurately than under pure-noise training (Appendix C.2).

Condition modalities and geometry. Unlike extending the action horizon, condition modalities exhibit structural redundancy and regularization rather than a simple monotonic ordering (geometric analysis in Appendix D).

7 Discussion

Across diverse manipulation benchmarks and physical bimanual tasks, rich conditioning over a compact target enables single-step flow matching without distillation. The velocity-loss profile diagnoses this structure from the Bayesian side; Gaussian concentration (Eq. (1)) explains it from the geometric side—a contractive transport map ($\sigma_1 < 1$ almost everywhere) renders the conditional target sub-Gaussian about its mean, while expansive image flows sit in the opposite regime. Yet compact is not collapsed: residual structure survives as sharp spectral fractures along physical decision boundaries, which flow matching causally encodes and neither direct regression nor pure-noise training resolves.

This shifts the scaling frontier from action decoding to condition construction. An agentic VLM actively enriches its own context through reasoning and interaction; because the resulting action target remains compact, decoding stays trivially cheap (Appendix G.1), rendering elaborate action-space diffusion machinery secondary to condition representation.

AI USE STATEMENT

In this work, we used generative AI tools for tasks with recommended disclosure: creating and editing software code (analysis and plotting scripts), creating and modifying scientific figures, and editing the manuscript to improve readability and grammar. We have not used generative AI tools for tasks with required disclosure; such tasks are not applicable to this work or were carried out by the authors. We have reviewed all AI-assisted work: AI-generated code was verified against experiment logs and outputs by the authors, and all AI-assisted figures and text were checked and approved by the authors. We take responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI.

REFERENCES

Michael S. Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *Journal of Machine Learning Research*, 26(209):1–80, 2025. URL https://jmlr.org/papers/v26/23-1605.html.

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, Mei Li, Kaixin Li, Zicheng Lin, Junyang Lin, Xuejing Liu, Jiawei Liu, Chenglong Liu, Yang Liu, Dayiheng Liu, Shixuan Liu, Dunjie Lu, Ruilin Luo, Chenxu Lv, Rui Men, Lingchen Meng, Xuancheng Ren, Xingzhang Ren, Sibo Song, Yuchong Sun, Jun Tang, Jianhong Tu, Jianqiang Wan, Peng Wang, Pengfei Wang, Qiuyue Wang, Yuxuan Wang, Tianbao Xie, Yiheng Xu, Haiyang Xu, Jin Xu, Zhibo Yang, Mingkun Yang, Jianxin Yang, An Yang, Bowen Yu, Fei Zhang, Hang Zhang, Xi Zhang, Bo Zheng, Humen Zhong, Jingren Zhou, Fan Zhou, Jing Zhou, Yuanzhi Zhu, and Ke Zhu. Qwen3-vl technical report, 2025. URL https://arxiv.org/abs/2511.21631.

Heli Ben-Hamu, Omri Puny, Itai Gat, Brian Karrer, Uriel Singer, and Yaron Lipman. D-flow: Differentiating through flows for controlled generation. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 3462–3483. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/ben-hamu24a.html.

Lucas Beyer, Andreas Steiner, André Susano Pinto, Alexander Kolesnikov, Xiao Wang, Daniel Salz, Maxim Neumann, Ibrahim Alabdulmohsin, Michael Tschannen, Emanuele Bugliarello, Thomas Unterthiner, Daniel Keysers, Skanda Koppula, Fangyu Liu, Adam Grycner, Alexey Gritsenko, Neil Houlsby, Manoj Kumar, Keran Rong, Julian Eisenschlos, Rishabh Kabra, Matthias Bauer, Matko Bošnjak, Xi Chen, Matthias Minderer, Paul Voigtlaender, Ioana Bica, Ivana Balazevic, Joan Puigcerver, Pinelopi Papalampidi, Olivier Henaff, Xi Xiong, Radu Soricut, Jeremiah Harmsen, and Xiaohua Zhai. PaliGemma: A versatile 3B VLM for transfer. arXiv preprint arXiv:2407.07726, 2024.

Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Robert Equi, Chelsea Finn, Niccolo Fusai, Manuel Y. Galliker, Dibya Ghosh, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Devin LeBlanc, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Allen Z. Ren, Lucy Xiaoyang Shi, Laura Smith, Jost Tobias Springenberg, Kyle Stachowicz, James Tanner, Quan Vuong, Homer Walke, Anna Walling, Haohuan Wang, Lili Yu, and Ury Zhilinsky. $\pi_{0.5}$: A vision-language-action model with open-world generalization. In *Proceedings of The 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pp. 17–40. PMLR, 2025a. URL https://proceedings.mlr.press/v305/black25a.html.

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Robert Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, Laura Smith, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. π_0: A vision-language-action flow model for general robot control. In *Robotics: Science and Systems*, 2025b. doi: 10.15607/rss.2025.xxi.010. URL https://www.roboticsproceedings.org/rss21/p010.html.

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

Siyi Chen, Huijie Zhang, Minzhe Guo, Yifu Lu, Peng Wang, and Qing Qu. Exploring low-dimensional subspace in diffusion models for controllable image editing. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 27340–27371. Curran Associates, Inc., 2024. doi: 10.52202/079017-0859. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/3018804d037cc101b73624f381bed0cb-Paper-Conference.pdf.

Zilong Chen, Chaorui Deng, Kunchang Li, Hongyi Yuan, and Haoqi Fan. Scaling properties of text conditioning in visual generation, 2026. URL https://arxiv.org/abs/2607.29679.

Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, 44(10–11):1684–1704, 2025. doi: 10.1177/02783649241273668.

Mingyang Deng, He Li, Tianhong Li, Yilun Du, and Kaiming He. Generative modeling via drifting, 2026. URL https://arxiv.org/abs/2602.04770.

Sander Dieleman. Generative modelling in latent space. Blog post, 2025. URL https://sander.ai/2025/04/15/latents.html.

Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis. In *ICML*, 2024.

Senyu Fei, Siyin Wang, Junhao Shi, Zihao Dai, Jikun Cai, Pengfang Qian, Li Ji, Xinzhe He, Shiduo Zhang, Zhaoye Fei, Jinlan Fu, Jingjing Gong, and Xipeng Qiu. LIBERO-Plus: In-depth robustness analysis of vision-language-action models. arXiv preprint arXiv:2510.13626, 2025.

Kevin Frans, Danijar Hafner, Sergey Levine, and Pieter Abbeel. One step diffusion via shortcut models. In *ICLR*, 2025.

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

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020.

Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, Kathrina Wu, Qin Lin, Junkun Yuan, Yanxin Long, Aladdin Wang, Andong Wang, Changlin Li, Duojun Huang, Fang Yang, Hao Tan, Hongmei Wang, Jacob Song, Jiawang Bai, Jianbing Wu, Jinbao Xue, Joey Wang, Kai Wang, Mengyang Liu, Pengyu Li, Shuai Li, Weiyan Wang, Wenqing Yu, Xinchi Deng, Yang Li, Yi Chen, Yutao Cui, Yuanbo Peng, Zhentao Yu, Zhiyu He, Zhiyong Xu, Zixiang Zhou, Zunnan Xu, Yangyu Tao, Qinglin Lu, Songtao Liu, Dax Zhou, Hongfa Wang, Yong Yang, Di Wang, Yuhong Liu, Jie Jiang, and Caesar Zhong. HunyuanVideo: A systematic framework for large video generative models. arXiv preprint arXiv:2412.03603, 2024.

Yann LeCun, Corinna Cortes, and Christopher J. C. Burges. The MNIST database of handwritten digits. Website, 1998. URL https://yann.lecun.com/exdb/mnist/.

Tianhong Li and Kaiming He. Back to basics: Let denoising generative models denoise. In *CVPR*, pp. 36115–36125, 2026.

Xuanlin Li, Kyle Hsu, Jiayuan Gu, Karl Pertsch, Oier Mees, Homer Rich Walke, Chuyuan Fu, Ishikaa Lunawat, Isabel Sieh, Sean Kirmani, Sergey Levine, Jiajun Wu, Chelsea Finn, Hao Su, Quan Vuong, and Ted Xiao. Evaluating real-world robot manipulation policies in simulation. *arXiv preprint arXiv:2405.05941*, 2024.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *ICLR*, 2023.

Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. LIBERO: Benchmarking knowledge transfer for lifelong robot learning. In *NeurIPS Datasets and Benchmarks*, pp. 44776–44791, 2023a.

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

Yiyang Lu, Susie Lu, Qiao Sun, Hanhong Zhao, Zhicheng Jiang, Xianbang Wang, Tianhong Li, Zhengyang Geng, and Kaiming He. One-step latent-free image generation with pixel mean flows, 2026. URL https://arxiv.org/abs/2601.22158.

Wuyang Luan, Junhui Li, Weiguang Zhao, Wenjian Zhang, Tieru Wu, and Rui Ma. SnapFlow: One-step action generation for flow-matching VLAs via progressive self-distillation. arXiv preprint arXiv:2604.05656, 2026.

Yuankai Luo, Woping Chen, Tong Liang, Baiqiao Wang, and Zhenguo Li. SimVLA: A simple VLA baseline for robotic manipulation. arXiv preprint arXiv:2602.18224, 2026.

Nanye Ma, Mark Goldstein, Michael S. Albergo, Nicholas M. Boffi, Eric Vanden-Eijnden, and Saining Xie. Sit: Exploring flow and diffusion-based generative models with scalable interpolant transformers, 2024. URL https://arxiv.org/abs/2401.08740.

Oier Mees, Lukas Hermann, Erick Rosete-Beas, and Wolfram Burgard. Calvin: A benchmark for language-conditioned policy learning for long-horizon robot manipulation tasks. *IEEE Robotics and Automation Letters (RA-L)*, 7(3):7327–7334, 2022.

Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Charles Xu, Jianlan Luo, Tobias Kreiman, You Liang Tan, Lawrence Yunliang Chen, Pannag Sanketi, Quan Vuong, Ted Xiao, Dorsa Sadigh, Chelsea Finn, and Sergey Levine. Octo: An open-source generalist robot policy. In *Robotics: Science and Systems*, Delft, Netherlands, 2024. doi: 10.15607/rss.2024.xx.090.

William Peebles and Saining Xie. Scalable diffusion models with transformers, 2023. URL https://arxiv.org/abs/2212.09748.

Trung X. Pham, Kang Zhang, Ji Woo Hong, and Chang D. Yoo. A hidden semantic bottleneck in conditional embeddings of diffusion transformers. In *ICLR*, 2026.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 2020.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pp. 10674–10685, 2022.

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

Shirin Shoushtari, Yi Wang, Xiao Shi, Salman Asif, and Ulugbek Kamilov. Eigenscore: Ood detection using posterior covariance in diffusion models. In C. Vondrick, B. Hariharan, C. Raffel, L. Pinto, D. Yang, and A. Faust (eds.), *International Conference on Learning Representations*, volume 2026, pp. 73590–73612, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/file/778b22d74907b01741bf1b6e82420b88-Paper-Conference.pdf.

Wenxuan Song, Jiayi Chen, Pengxiang Ding, Yuxin Huang, Han Zhao, Donglin Wang, and Haoang Li. CEED-VLA: Consistency vision-language-action model with early-exit decoding. arXiv preprint arXiv:2506.13725, 2025.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*, 2021.

Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *ICML*, 2023.

Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling: Scalable image generation via next-scale prediction. In *NeurIPS*, pp. 84839–84865, 2024.

Jingfeng Yao, Bin Yang, and Xinggang Wang. Reconstruction vs. generation: Taming optimization dilemma in latent diffusion models. In *CVPR*, pp. 15703–15712, 2025.

Jinhui Ye, Ning Gao, Senqiao Yang, Jinliang Zheng, Zixuan Wang, Yuxin Chen, Pengguang Chen, Yilun Chen, Shu Liu, and Jiaya Jia. Starvla-α: Reducing complexity in vision-language-action systems. In *European Conference on Computer Vision (ECCV)*, 2026.

Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Frédo Durand, William T. Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *CVPR*, pp. 6613–6623, 2024.

Lijun Yu, Yong Cheng, Kihyuk Sohn, José Lezama, Han Zhang, Huiwen Chang, Alexander G. Hauptmann, Ming-Hsuan Yang, Yuan Hao, Irfan Essa, and Lu Jiang. MAGVIT: Masked generative video transformer. In *CVPR*, pp. 10459–10469, 2023.

Lijun Yu, José Lezama, Nitesh Bharadwaj Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G. Hauptmann, Boqing Gong, Ming-Hsuan Yang, Irfan Essa, David Ross, and Lu Jiang. Language model beats diffusion: Tokenizer is key to visual generation. In *ICLR*, 2024.

Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *ICCV*, pp. 11941–11952, 2023.

Shiduo Zhang, Yue Wang, Haonan Chang, Hang Zhao, Yicheng Liu, Vitor Guizilini, Andreea Bobu, Andrew Wagenmaker, Anshul Dixit, Chuan Yu, Dhruv Shah, and Max Simchowitz. Post-training for robotics foundation models dataset and challenge. RSS 2026 Workshop & Challenge, 2026. URL https://posttraining-for-robotics.github.io.

Jinliang Zheng, Jianxiong Li, Zhihao Wang, Dongxiu Liu, Xirui Kang, Yuchun Feng, Yinan Zheng, Jiayin Zou, Yilun Chen, Jia Zeng, Ya-Qin Zhang, Jiangmiao Pang, Jingjing Liu, Tai Wang, and Xianyuan Zhan. X-vla: Soft-prompted transformer as scalable cross-embodiment vision-language-action model, 2025. URL https://arxiv.org/abs/2510.10274.

Xueyang Zhou, Yangming Xu, Guiyao Tie, Yongchao Chen, Guowen Zhang, Duanfeng Chu, Pan Zhou, and Lichao Sun. LIBERO-PRO: Towards robust and fair evaluation of vision-language-action models beyond memorization. arXiv preprint arXiv:2510.03827, 2026.

A PROOF OF EQ. (1)

Theorem 1 (Gaussian transport concentration)**.** *Fix a condition c, let $z \sim \mathcal{N}(0, I_d)$, and define $x_1 = \phi_1(z; c)$, where $\phi_1(\cdot; c) \in W^{1,\infty}_{\mathrm{loc}}(\mathbb{R}^d; \mathbb{R}^d)$ has Jacobian $J(z; c) = \nabla_z \phi_1(z; c)$. Let*

$$L(c) \triangleq \operatorname*{ess\,sup}_z \|J(z; c)\|_2 < \infty.$$

Then every unit projection $u^\top x_1$, $u \in \mathbb{S}^{d-1}$, is $L(c)^2$-sub-Gaussian:

$$\mathbb{E}_z\Big[\exp\Big(\lambda\, u^\top \big(x_1 - \mathbb{E}_z[x_1]\big)\Big)\Big] \le \exp\left(\frac{\lambda^2 L(c)^2}{2}\right), \qquad \forall\, \lambda \in \mathbb{R}. \tag{2}$$

Moreover,

$$\mathbb{P}_z\left(\|x_1 - \mathbb{E}_z[x_1]\|_2 \ge \sqrt{\mathbb{E}_z\|J(z; c)\|_F^2} + \epsilon\right) \le \exp\left(-\frac{\epsilon^2}{2L(c)^2}\right), \qquad \forall\, \epsilon > 0. \tag{3}$$

Since $\mathbb{E}_z\|J(z; c)\|_F^2 \le d\, L(c)^2$, the characteristic radius is further upper bounded by $L(c)\sqrt{d}$.

Proof. Since $\mathbb{R}^d$ is convex and $\|J\|_2 \le L(c)$ almost everywhere, $\phi_1(\cdot; c)$ is globally $L(c)$-Lipschitz. In particular, all expectations below are finite.

Projections. For $u \in \mathbb{S}^{d-1}$, let $g(z) = u^\top \phi_1(z; c)$. Then

$$\|\nabla g(z)\|_2 = \|J(z; c)^\top u\|_2 \le L(c) \quad \text{almost everywhere.}$$

The standard Gaussian measure γ satisfies the logarithmic Sobolev inequality

$$\mathrm{Ent}_\gamma(f^2) \le 2\, \mathbb{E}_\gamma\big[\|\nabla f\|_2^2\big].$$

Applying it with $f = e^{\lambda g/2}$ and using the Herbst argument yields

$$\mathbb{E}_z\big[\exp\big(\lambda(g - \mathbb{E}_z[g])\big)\big] \le \exp\left(\frac{\lambda^2 L(c)^2}{2}\right) \tag{4}$$

for $\lambda > 0$; applying the same argument to $-g$ gives the result for all $\lambda \in \mathbb{R}$.

Norm. Define

$$F(z) = \big\|\phi_1(z; c) - \mathbb{E}_z[\phi_1(z; c)]\big\|_2.$$

Since $\phi_1(\cdot; c)$ is $L(c)$-Lipschitz, the reverse triangle inequality shows that F is $L(c)$-Lipschitz. Hence Gaussian concentration gives

$$\mathbb{P}_z(F \ge \mathbb{E}_z[F] + \epsilon) \le \exp\left(-\frac{\epsilon^2}{2L(c)^2}\right). \tag{5}$$

Next, since $\mathbb{E}_z[F^2] = \sum_{i=1}^d \mathrm{Var}_z\left(\phi_1^{(i)}\right)$, the Gaussian Poincaré inequality applied coordinate-wise gives

$$\mathbb{E}_z[F^2] = \sum_{i=1}^d \mathrm{Var}_z\Big(\phi_1^{(i)}(z; c)\Big) \le \sum_{i=1}^d \mathbb{E}_z\Big[\big\|\nabla_z \phi_1^{(i)}(z; c)\big\|_2^2\Big] = \mathbb{E}_z\|J(z; c)\|_F^2. \tag{6}$$

Therefore, by Jensen's inequality,

$$\mathbb{E}_z[F] \le \sqrt{\mathbb{E}_z[F^2]} \le \sqrt{\mathbb{E}_z\|J(z; c)\|_F^2}.$$

Consequently,

$$\left\{F \ge \sqrt{\mathbb{E}_z\|J(z; c)\|_F^2} + \epsilon\right\} \subseteq \left\{F \ge \mathbb{E}_z[F] + \epsilon\right\},$$

and Eq. (1) follows from Eq. (5).

Radius bound. Since $J(z; c)$ is $d \times d$, $\|J(z; c)\|_F^2 \le d\, \|J(z; c)\|_2^2 \le d\, L(c)^2$ almost everywhere, which yields the stated coarse radius bound. □

B TIME COORDINATES AND TRAINING SCHEDULES

The paper uses the flow-matching coordinate t, with $t = 0$ denoting the noise endpoint and $t = 1$ denoting clean data. OpenPI uses the opposite coordinate $t_{\text{op}} = 1 - t$, where $t_{\text{op}} = 1$ is pure noise. Its default base distribution $t_{\text{op}} \sim \text{Beta}(1.5, 1)$ therefore becomes $t \sim \text{Beta}(1, 1.5)$ in the paper coordinate. For an OpenPI-coordinate sample, the high-noise shift is $t_{\text{op,shifted}} = \alpha t_{\text{op}}/(1 + (\alpha - 1)t_{\text{op}})$. Larger α moves probability mass toward $t_{\text{op}} = 1$, equivalently $t = 0$, while preserving the velocity-prediction target. This is the same schedule as in Section 5.1 after changing coordinates.

All diagnostics are plotted in the paper coordinate (OpenPI logs converted as $t = 1 - t_{\text{op}}$).

Velocity-MSE protocol. All velocity-MSE profiles (Figs. 4, 5a, and 9) are computed by the same fixed-time diagnostic. For each checkpoint we sample a fixed set of evaluation contexts and, for each context, draw a matched Gaussian $x_0 \sim \mathcal{N}(0, I)$; the matched-noise anchor protocol draws a single long noise sequence and slices prefixes of length H, so every horizon sees identical noise. We then evaluate the model at 25 paper times $t \in [0.001, 0.999]$ and regress the predicted velocity against the analytic target $u = x_1 - x_0$, reporting the mean squared error over the seven physical action dimensions (the 32-dimensional action tensor is sliced to its first seven coordinates; the padding coordinates never enter any loss or metric). Each profile aggregates 64 context–noise pairs.

Table 6: **Direct learned action-start controls on standard LIBERO** (success rates %). All rows use the Tiny H10 policy with one-step decoding (500 episodes per suite). Differences (±) are relative to the pure-noise 1-step Tiny baseline (Table 1).

Learned Start Init	Spatial	Object	Goal	Long
Zero	92.8 -3.2	99.0 +1.2	59.4 -36.8	68.2 -4.6
Normal, std = 1	95.4 (-0.6)	98.6 +0.8	59.8 -36.4	69.4 -3.4

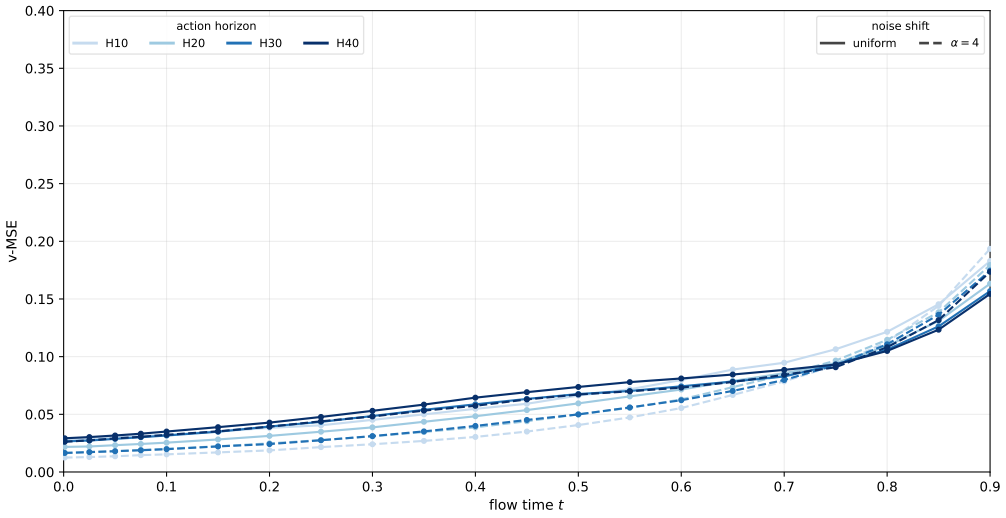

Figure 9: Full horizon sweep underlying Fig. 5a, including the H10 curves. Color denotes the horizon (H10 light to H40 dark); linestyle denotes the training-time distribution (uniform solid, α=4 dashed).

C TRANSPORT-GEOMETRY DIAGNOSTICS

This appendix specifies how every σ_1 quantity in Section 6 is computed. Unless stated otherwise, all estimators share the same object: the noise-to-action map $\phi_1^\theta(\cdot\,; c)$ and its Jacobian $J(x_0; c) = \nabla_{x_0}\phi_1^\theta(x_0; c)$.

C.1 PROTOCOLS

Endpoint Jacobian. For models that natively parameterize one-step transport (MeanFlow, pMF) we differentiate the single forward map directly. For continuous flow models (π_0, $\pi_{0.5}$, StarVLA, X-VLA, and our LIBERO checkpoints) the trained vector field $v_\theta(x, t, c)$ is integrated along the paper-time path from $t = 0$ to $t = 1$ with 10-step Euler integration, and ϕ_1^θ denotes the composed

map; J is obtained by reverse-mode autodifferentiation through the composition. At each evaluation point we form the Jacobian and take its largest singular value, either by a direct SVD for small action dimensions or by power iteration on $J^\top J$.

Transition screening. Gripper transitions are screened from held-out rollouts of the H20 uniform-schedule policy: around each recorded gripper close we draw a dense candidate set of 444 noise sources, evaluate the ϕ_1^θ spectrum of each, and retain high-σ_1 states with an observable gripper switch at chunk step m; greedy no-replacement selection across episodes yields 32 events.

Steering scans. Each event state is re-decoded under shifted noise $x_0 + \epsilon v_1$, with ϵ on the symmetric grid $\{0, \pm 0.02, \pm 0.05, \pm 0.1, \pm 0.2, \pm 0.5\}$; the switch step m is read off the decoded gripper trace, and an event counts as switched when $m(+\epsilon) \neq m(-\epsilon)$ at some paired magnitude (Fig. 10a). The main text reports only $\epsilon \leq 0.1$, where shifted probes remain in the typical Gaussian region.

Pure-noise steering control. The pure-noise policy is natively one-step and has no multi-step flow to compose, so its control probe targets its own one-step map: the H20 pure-noise checkpoint is screened by the same high-σ_1 procedure on that map and re-decoded under the same ϵ grid and switch statistic as above.

Landscape cross-sections. For one transition (episode 18, step 5116; H20 uniform-schedule policy) we evaluate the fp32 flow map ϕ_1^θ (10-step Euler composition) on a 70×81 grid in the (v_1, v_2) plane, producing the landscape of Fig. 8b; a circular orbit of radius 0.06 and a 1D decode lattice along a at b=0 produce the orbit decodes (Fig. 8c), the σ cross-section, and the quantized switch step $m(a)$ (Fig. 10c,d).

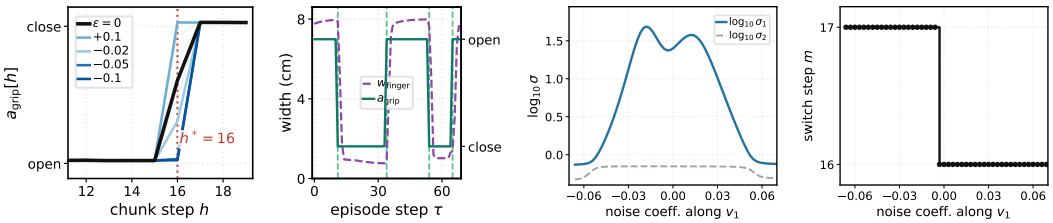

(a) causal steering (v_1) (b) actuation trace (c) core cross-section (d) decoded switch step

Figure 10: **Steering scans and landscape cross-sections.** (a) Commanded gripper action a_{grip} along chunk steps h under initial noise shifts $x_0 + \epsilon v_1$. (b) Commanded gripper action a_{grip} and physical finger width w_{finger} confirm that fractures coincide with physical contact transitions. (c) Cross-section at b=0 of the ϕ_1^θ landscape in Fig. 8: σ_1 (blue) forms two lanes flanking a saddle, while σ_2 (gray dashed) stays below 0.7. (d) Decoded switch step $m(a)$ at b=0: a single jump inside the saddle region.

Episode-time traces. Figure 8a follows the H20 uniform-schedule policy along one held-out rollout segment, sampling an observation every 8 frames around its recorded gripper boundaries (green markers). At each observation c_τ, conditioned on the recorded frame, σ_1 of the complete-map Jacobian $\nabla_z \phi_1^\theta(\,\cdot\,; c_\tau)$—the same 10-step Euler map as in the landscape protocol above—is estimated matrix-free by block power iteration (block rank 2, three subspace iterations) at $K = 4$ fixed Gaussian noise sources, and the mean $\log_{10}\sigma_1$ over sources is plotted. The trace stays flat and contractive between boundaries, climbs across σ_1=1 ahead of each boundary, and drops back immediately after; the commanded gripper action and physical finger width around such transitions are shown in Fig. 10b.

Noise-plane heatmaps. Figure 7a evaluates a latent MeanFlow model on the plane $z_0 + av_1 + bv_2$ around a fixed anchor z_0; σ_1 is evaluated on a 17×17 grid by per-point autodiff, and the orbit $z(\alpha)$ is the circle of radius 0.15 in this plane. The cyan contour marks the p95 ridge of the grid spectrum.

Population distributions. The histograms of Figs. 6 and 13 are built by sampling evaluation contexts $c = (o_t, l, s_t)$ from held-out demonstrations together with i.i.d. Gaussian noise x_0, evaluating $\sigma_1(x_0; c)$ at each pair (1088 pairs per condition cell), and plotting the density in $\log_{10}\sigma_1$. Fig. 6 uses the endpoint Jacobian; Fig. 13 uses ϕ_1^θ (10-step Euler) of the H10 α=4 policy.

Schedule specificity. Figure 11 repeats the per-step diagnostic on the same episode under the three training schedules, now with the one-step generator map $T_c(z) = z + v_\theta(z, t{=}0; c)$ ($K{=}16$ fixed noise anchors, full SVD) so that the natively one-step pure-noise policy enters on equal footing. The quiet-between-transitions structure is specific to the flow map: the pure-noise trace stays noisy throughout. The same specificity holds under intervention: at matched small shifts ($\epsilon \leq 0.1$) the v_1-steering probe of Section 6.2 changes m in only 6.5% of screened pure-noise events (18.3% at $\epsilon \leq 0.2$; 34.4% over the full grid), whereas ϕ_1^θ responds reliably already at $\epsilon \leq 0.1$.

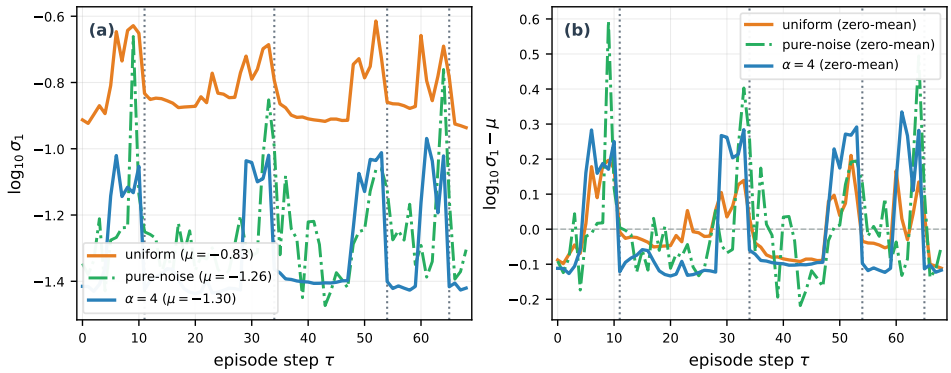

Figure 11: Episode geometry across training schedules. (a) Per-step $\log_{10}\sigma_1$ of the same episode under uniform, pure-noise, and $\alpha{=}4$ training (means -0.83, -1.26, -1.30). (b) Zero-mean relative sensitivities around the same boundaries.

C.2 Boundary Timing of the Deployed One-Step Map

Section 6.2 probes the composed flow map ϕ_1^θ (10-step Euler), while deployment decodes by a single vector-field evaluation $T_c(z) = z + v_\theta(z, t{=}0; c)$. The deployed map still carries the boundary's timing: at the same 32 screened anchors (H20 uniform), h^* falls within two chunk steps of the recorded gripper transition for 88% of anchors, whereas the pure-noise map's h^* peaks systematically three steps early (28%; Fig. 12).

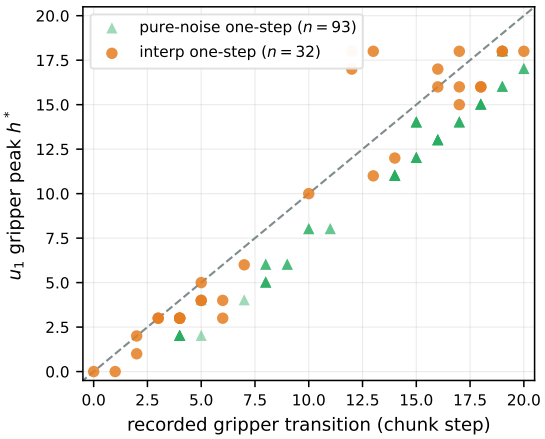

Figure 12: **The deployed one-step map retains boundary timing.** Gripper-channel peak h^* of the one-step map's leading output direction vs. the recorded gripper transition step, at the 33 screened anchors (H20 uniform): the interpolation-trained policy tracks the physical timing, while the pure-noise control (screened at its own anchors) peaks systematically early. Dashed: exact match.

C.3 Output Normalization and Cross-Domain Spectra

The spectral norm $\sigma_1(J)$ is coordinate dependent, and the two domains in Fig. 6 use different output conventions: action coordinates are per-coordinate standardized with training-set mean and standard deviation (unit marginal variance), whereas image models use the conventional pixel scaling $x_i \in [-1, 1]$. The following bound shows this mismatch cannot explain the cross-domain gap: variance-matching the image coordinates would only enlarge the image spectra.

Proposition 1 (Coordinate-wise standardization). *Let* $D = \mathrm{diag}(s_1, \ldots, s_d)$, $s_i > 0$, *and* $\widetilde{x} = D^{-1}(x - \mu)$ *where* $x = \phi_1^\theta(z; c)$. *Translation does not affect the Jacobian, so* $J_{\widetilde{x}} = D^{-1}J$, *and*

$$\frac{\sigma_1(J)}{s_{\max}} \;\le\; \sigma_1(D^{-1}J) \;\le\; \frac{\sigma_1(J)}{s_{\min}}, \qquad s_{\min} = \min_i s_i,\; s_{\max} = \max_i s_i.$$

Proof. For any unit vector v, $\|D^{-1}Jv\|_2 \ge \sigma_{\min}(D^{-1})\,\|Jv\|_2 = \|Jv\|_2/s_{\max}$ and $\|D^{-1}Jv\|_2 \le \|D^{-1}\|_2\,\|Jv\|_2 = \|Jv\|_2/s_{\min}$; taking the supremum over $\|v\|_2 = 1$ gives both bounds. □

Every image coordinate is supported on $[-1, 1]$, hence its marginal standard deviation satisfies $s_i^{\mathrm{img}} \le 1$. Standardizing image outputs to unit marginal variance therefore gives $\sigma_1(J_{\mathrm{img,std}}) \ge \sigma_1(J_{\mathrm{img}})/s_{\max}^{\mathrm{img}} \ge \sigma_1(J_{\mathrm{img}})$: the reported comparison is conservative with respect to this scale mismatch. This statement concerns only coordinate-wise scaling—σ_1 is not claimed invariant under general linear reparameterizations such as full covariance whitening—and all measurements are taken in each model's native training coordinates.

D Condition Ablations and Closed-Loop Proprioceptive Failure

We ablate individual condition components of the H10 $\alpha = 4$ one-step policy (Table 7, 500 episodes per suite), and the degradations are highly non-uniform. Visual ablation leaves basic reaching intact on simpler suites—the Object suite swaps only the target object at fixed layouts, so it remains solvable from proprioception and task context—but collapses long-horizon tasks (Long drops to 31.8% and 53.2%). Language ablation causes an isolated semantic collapse on Goal (8.8%), where scenes are identical across tasks and the prompt is the sole grounding signal. Proprioceptive ablation causes widespread prehensile failure (Object 18.2%, Long 2.0%) while Goal largely survives (44.0%).

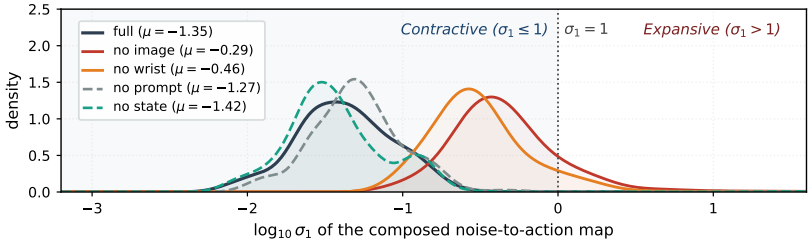

Figure 13: **Transport geometry under condition ablations.** Distribution of $\log_{10}\sigma_1$ of the noise-to-action map for the H10 α=4 policy (Table 7): removing the image or the wrist camera pushes mass across the contractive boundary σ_1=1 (dotted), i.e., a genuinely broadened conditional distribution; removing the prompt or the state leaves the geometry essentially unchanged—prompt deletion in particular induces categorical semantic ambiguity rather than continuous geometric dispersion.

Table 7: Condition ablations on standard LIBERO (H10 α=4 one-step tiny policy, success rates %, 500 episodes per suite).

Condition	Spatial	Object	Goal	Long
Full	96.4	99.6	96.8	85.2
No image	82.2	96.0	85.4	31.8
No prompt	82.0	97.2	8.8	55.2
No wrist	92.2	96.4	88.0	53.2
No state	33.8	18.2	44.0	2.0

Closed-loop failure of the no-state policy. The no-state ablation leaves the transport geometry essentially unchanged—σ_1 even contracts slightly (Fig. 13)—and the policy tracks recorded demonstrations as accurately as the full policy in open loop (per-step end-effector error ratio 0.93–1.03): LIBERO's fixed scenes make the arm state largely estimable from vision. What fails is closed-loop control, and a 2×2 same-root intervention (Table 8) shows the state channel itself is causal: restarting from exact environment snapshots of failed no-state rollouts—taken just before or at the grasp—the full policy recovers 94.4% and 91.7% of them respectively, the no-state policy itself only 5.6%. Our tentative reading: the state channel matters less as extra information than as making the condition structure explicit and easy to learn, which is what generalizes off the demonstrated trajectory.

Table 8: **Same-root intervention** (success rates %). Paired snapshots of successful full-policy and failed no-state rollouts (LIBERO Object) are taken at matched timesteps before the grasp (approach) and at the grasp (event), then continued under both policies; full/no-state columns denote which rollout supplied the snapshot. 36 continuations per cell (18 snapshots $\times$ 2 noise sources).

	Approach roots		Event roots	
Policy	full	no-state	full	no-state
Full	100.0	94.4	100.0	91.7
No-state	38.9	5.6	50.0	5.6

E Image-to-Text MNIST Details

The condition is a 128×128 image of a 4×4 grid of MNIST digits (LeCun et al., 1998), and the target is the corresponding 16-token digit sequence, read in row-major order. A CNN encoder embeds the grid image into condition tokens. The target token embeddings $x_1 \in \mathbb{R}^{16\times256}$ follow the standard linear interpolation $x_t = tx_1 + (1-t)x_0$ with $x_0 \sim \mathcal{N}(0, I)$. The noised target tokens and the condition tokens are concatenated and fed into a bidirectional Transformer decoder without causal masking, which predicts the velocity $v_\theta(x_t, t, c)$.

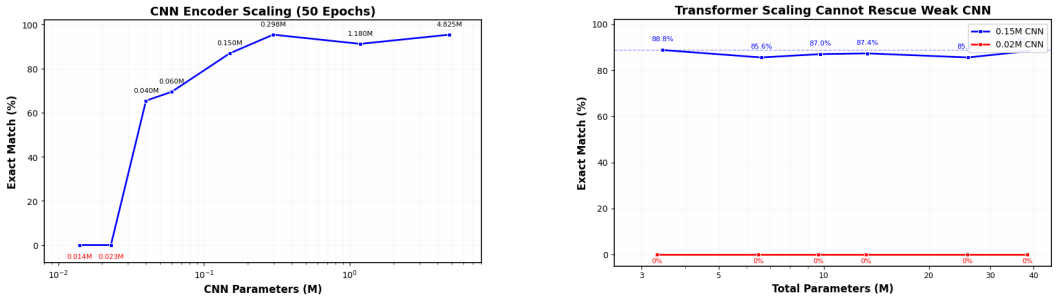

(a) CNN encoder scale.

(b) Decoder scale compensation.

Figure 14: Two ablations: (a) varying the CNN encoder scale shows that weaker encoders degrade one-step performance; (b) scaling up the Transformer decoder while keeping a weak encoder does not compensate, confirming that the bottleneck lies in the condition-target structure, not model capacity.

The image-to-text MNIST dataset contains 3,750 training examples and 625 test examples. In addition to the flow-matching loss, we use a cross-entropy term on the clean-target prediction, $\mathcal{L} = \mathcal{L}_{\text{FM}} + \lambda\mathcal{L}_{\text{CE}}, \lambda = 2$, where $x_{\text{pred}} = x_t + (1-t)v_{\text{pred}}$ is used to evaluate $\mathcal{L}_{\text{CE}}$. Table 9 gives the shared training configuration behind the MNIST schedule and encoder/decoder controls in the main text.

Table 9: Image-to-text MNIST training hyperparameters.

Parameter	Value
Image size	128×128
Maximum text length	20 tokens
Vocabulary size	13
d_{model} / layers / heads	256 / 4 / 8
Feedforward dimension	1024
Dropout	0.1

Parameter	Value
Batch size	128
Learning rate	10^{-3}
Weight decay	0.01
Optimizer / scheduler	AdamW / cosine
Warmup / total epochs	5 / 50
Min LR / grad clip	10^{-6} / 1.0
Prediction mode	v-prediction

F Robot Policy Architecture and Training

Tiny ablations use a four-layer SigLIP slice and a four-layer Gemma stack. The full-encoder validation restores all 27 SigLIP layers and the first four pretrained PaliGemma layers while retaining the same shallow action head. That head concatenates condition tokens, state tokens, and action–time tokens, then applies a pre-LayerNorm bidirectional Transformer and projects action tokens to velocities. Tables 11 and 12 specify the reported configurations.

Table 10: Parameter breakdown for Tiny and Full policy architectures. Active parameters exclude the frozen vocabulary embedding table.

Component	**Tiny (108.7M active)**		**Full (887.6M active)**	
	Parameters	Size (MB)	Parameters	Size (MB)
Vision Encoder (SigLIP)	62.52M	238.5	414.80M	1582.3
Multimodal LLM (Gemma backbone)	14.95M	57.0	440.42M	1679.5
Action Flow Head (+ Proj, Time MLP)	**31.23M**	118.3	**32.41M**	123.1
Total (incl. Text Embedding)	240.40M	917.1	1414.30M	5395.2

Table 11: Architecture for the reported LIBERO runs. Horizon controls change only `action_horizon`; H10 denotes horizon 10.

Component	Parameter	Tiny	Full encoder
Gemma/VLM	variant	`gemma_tiny`	`gemma_2b`
	used layers	4	4
	width	512	2048
	mlp_dim	2048	16384
	num_heads	8	8
	num_kv_heads	1	1
	head_dim	64	256
Vision	encoder	SigLIP-So400m	SigLIP-So400m
	used layers	4	all 27
Action head	mode / depth	SimVLA / 4	SimVLA / 4
	width	768	768
	heads	12	12
	mlp_dim	3072	3072
Action	action_dim	32	32
	action_loss_dim	7	7
	action_horizon	10	10

Table 12: Training hyperparameters for the reported LIBERO runs.

Parameter	Tiny ablations	Full encoder
Optimizer	AdamW	AdamW
$\beta_1, \beta_2, \epsilon$	0.9, 0.95, 10^{-8}	0.9, 0.95, 10^{-8}
Weight decay / gradient clip	10^{-10} / 1.0	10^{-10} / 1.0
LR schedule	cosine decay	cosine decay
Warmup / decay steps	1k / 50k	1k / 150k
Peak / final LR	10^{-4} / 10^{-5}	2×10^{-4} / 2×10^{-5}
Backbone LR multiplier	0.1	0.1
EMA decay	0.99	0.99
Batch size / steps	64 / 50k	256 / 150k



Figure 15: Overview of real-world bimanual manipulation tasks on the DualYam robot platform.

G Evaluation Protocol and Controls

A policy query produces an action chunk of horizon H. We separately report the number of Euler solver steps (`num_flow_steps`, abbreviated fs) and the number of environment actions executed before the next policy query (`replan_steps`, abbreviated r). Unless stated otherwise, the main results use H10 and execute the entire predicted chunk ($r = 10$). This distinction matters: $fs = 1$ specifies one flow evaluation per query, not an open-loop rollout length.

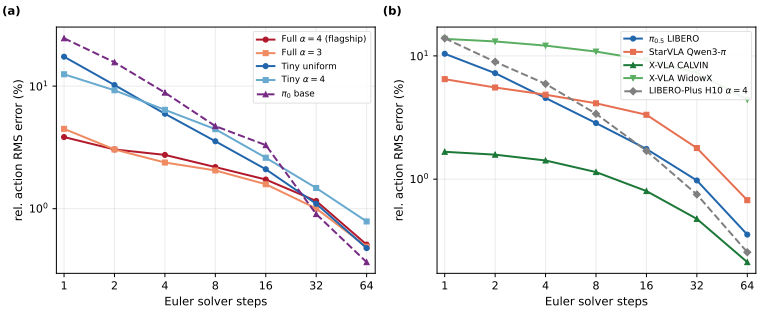

Figure 16: Solver-step convergence audit. Relative action RMS error of Euler integration with $fs \in \{1, 2, 4, 8, 16, 32, 64\}$ steps against a 128-step Euler reference, measured on fixed observations. (a) Our LIBERO checkpoints; (b) third-party checkpoints and a LIBERO-Plus recipe. The error drops below 1% by 64 steps for all but X-VLA WidowX, while the one-step error remains small (1.7–19%), consistent with the strong-condition signature of Fig. 4.

Table 13 gives the H10 replanning sensitivity; all comparisons keep fs, r, and horizon separate.

Table 13: H10 replanning sensitivity on LIBERO-Long. Values are success rates (%).

Recipe	Flow steps	$r = 1$	$r = 2$	$r = 5$	$r = 10$
Uniform	1	32.0	38.4	72.8	70.2
Uniform	10	55.0	58.0	80.2	80.8
Noise shift ($\alpha = 3$)	1	24.0	33.2	79.2	78.0
Noise shift ($\alpha = 3$)	10	15.6	25.4	52.4	52.4
Noise shift ($\alpha = 4$)	1	33.0	45.0	86.2	85.2
Noise shift ($\alpha = 4$)	10	18.0	29.0	66.8	63.4
Noise shift ($\alpha = 8$)	1	29.4	39.8	81.4	78.0
Noise shift ($\alpha = 8$)	10	3.4	7.0	22.6	22.6
Pure noise	1	23.2	33.2	73.6	72.8

G.1 Inference Latency Decomposition

Figure 17 decomposes per-query model-forward latency (NVIDIA RTX 5090, batch size 1, mean of 50 warmed calls on synthetic observations with the deployed token counts; excludes input/output transforms, networking, and the simulator) into the condition prefix and the action decoder, at each architecture's multi-step baseline and at one-step decoding. Single-stepping cuts total latency by 1.5–3.4× across all five architectures. The condition prefix then dominates the remainder for our policies (52% and 62% of the Tiny and Full forwards) and accounts for 83% of the official π_0 forward—the systems-level counterpart of the condition-dominance picture: once decoding is single-step, the remaining cost is condition construction, not action generation. Our two columns are profiled from the deployed JAX path (checkpoints of Tables 1 and 2); the π_0 column is profiled on the official DROID checkpoint, which shares its architecture and horizon with the LIBERO one; the X-VLA and StarVLA columns reuse our earlier profiling under the same protocol.

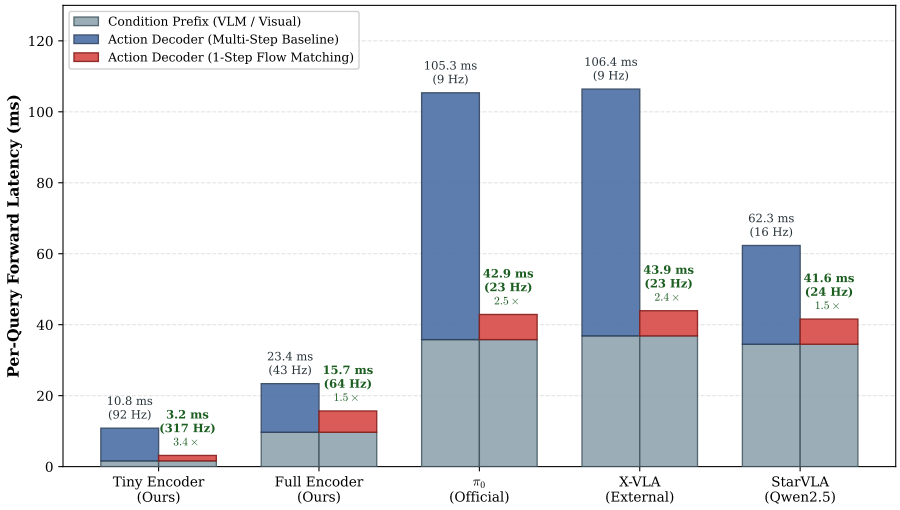

Figure 17: **Inference latency decomposition** (RTX 5090, batch 1, model forward). Condition prefix vs. action decoder at the multi-step baseline (10-step; 4-step for StarVLA) and at one-step decoding. Single-stepping cuts total per-query latency 1.5–3.4×, after which the condition prefix dominates the remainder.

H LIBERO-PLUS FULL-CONDITION SWEEP

We evaluate the tiny policy directly on LIBERO-Plus (Fei et al., 2025) across horizons and noise schedules (2500 episodes/suite). As shown in Fig. 18a, across 18 comparable non-pure recipes, one-step decoding consistently matches or exceeds ten-step decoding (+5.4 pp mean margin), confirming that the high-noise one-step benefit generalizes under distribution shifts.

LIBERO-Pro robustness probe. We test the full-encoder standard-LIBERO checkpoint directly on LIBERO-Pro perturbations (Zhou et al., 2026) without task-specific tuning (500 episodes/cell, Fig. 18b). One-step decoding averages 44.2% across all cells, while ten-step averages 43.5% (15 of 16 cells within 5 points; per-perturbation one-step means: Swap 2.6, Object 75.5, Language 82.6, Task 16.2), confirming that one-step decoding is robust to perturbations without sacrificing generalization.

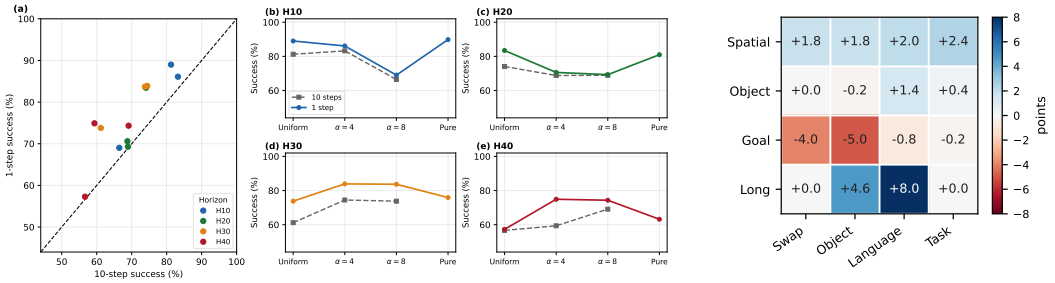

Figure 18: **One-step parity under distribution shift.** (a) LIBERO-Plus full-condition sweep across horizons and schedules: all 18 comparable non-pure recipes lie on or above the diagonal (+5.4 pp mean margin). (b) LIBERO-Pro robustness probe: success-rate difference between 1-step and 10-step decoding across the 16 perturbation cells.

Table 14 provides the exact aggregate values. Each recipe is evaluated across the four suites with around 2,500 episodes per suite. Values are means over the four suites; blank cells correspond to recipes not trained in the sweep.

Table 14: LIBERO-Plus full-condition one-step sweep. U denotes uniform; numbered columns denote α. Values are four-suite mean success rates (%).

Horizon	U	1	2	3	4	8	Pure
H10	89.0	92.9	94.3	80.7	86.1	69.0	89.8
H20	83.4	68.2	49.1	63.6	70.6	69.3	81.0
H30	73.8	–	–	–	83.9	83.7	75.9
H40	57.3	–	–	–	74.9	74.3	63.2

I ACTION-HEAD POSITIONAL CONTROLS

The deeper-head ablations use the same tiny-model setup as the main LIBERO controls; the per-layer position variant adds a learned absolute position embedding inside each action-decoder block. These are preliminary probes, not an extensive study: every variant is trained once under the shared hyperparameters of Table 12 with no per-variant tuning, so the numbers are directional evidence rather than optimized comparisons. They do not support a blanket claim that decoder design is irrelevant: deeper heads can preserve one-step quality, but naive scaling makes multi-step inference fragile and yields no consistent gain.

Table 15: Depth-8 action-head controls (1-step / 10-step success %). Top: per-suite under H10 α=4. Bottom: LIBERO-Long vs. train-time α.

Variant	Train α	Spatial	Object	Goal	Long
Vanilla depth8	4	95.6 / 70.0	99.0 / 53.6	97.0 / 69.6	80.2 / 2.6
Per-layer position	4	97.4 / 84.6	98.8 / 83.2	95.4 / 73.4	60.2 / 9.6
RoPE	4	97.0 / 69.6	99.2 / 57.8	94.4 / 65.8	80.0 / 8.8
Shallow baseline	4	–	–	–	85.2 / 63.4
Depth8 RoPE	3	–	–	–	74.6 / 41.2
Depth8 RoPE	8	–	–	–	70.2 / 25.8