Flow Model Post-Training as Representation Distribution Matching

Lan Feng[1,2] Peter Karkus[2] Maximilian Igl[2] Yuxiao Chen[2]
Daniel Dauner[3] Alexander Naumann[2] Jef Packer[2] Wenhao Ding[2]
Xiangyu Chen[2] Alexandre Alahi[1] Boris Ivanovic[2] Marco Pavone[2,4]

[1]EPFL [2]NVIDIA [3]KE:SAI [4]Stanford University

Project page: `https://alan-lanfeng.github.io/FlowRDM-website/`

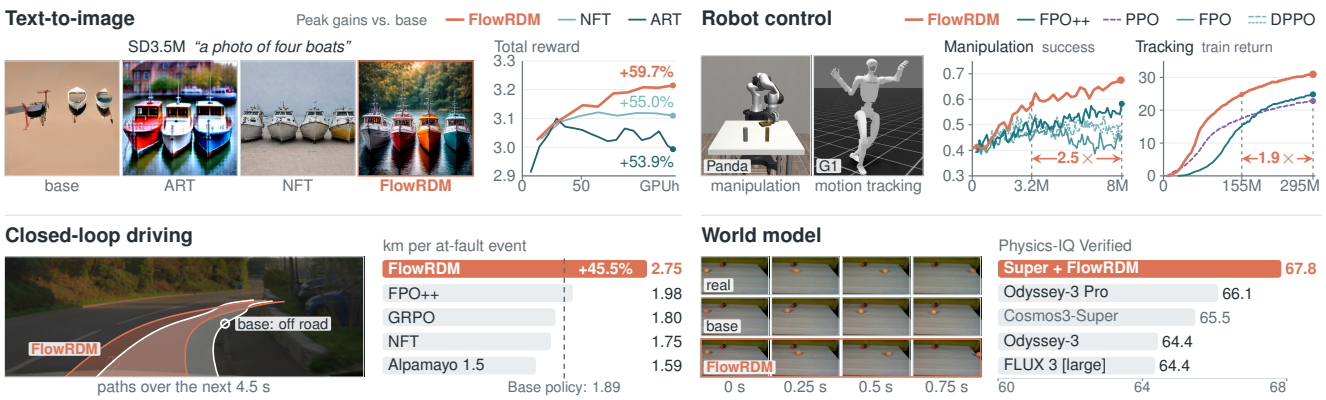

Figure 1. **FlowRDM improves visual generation and embodied control.** Reward-weighted representation distribution matching transfers across diverse flow models, matching or surpassing state-of-the-art methods. Coral denotes FlowRDM; percentages give gains over the base models, and arrows show environment-step ratios to reach FPO++'s final performance (Appendix F).

Abstract

Flow models are often post-trained with non-differentiable rewards. Most methods weight samples independently or move a group along one mean direction, ignoring the relations among samples. We propose FlowRDM, a unified framework that formulates flow-model post-training as representation distribution matching (RDM). It matches samples to their reward-weighted distribution in feature space, letting high-reward samples guide updates without reward gradients. In embodied control, it compares actions sampled at the same visited state. Across diverse tasks and base models in visual generation and embodied control, it matches or surpasses state-of-the-art methods. On SD3.5-Medium and FLUX.1-dev, it raises total reward by 59.7% and 14.1% over the respective base models, outperforming DiffusionART and DiffusionNFT on total reward and two reward models excluded from training. Post-trained on simulated videos alone, Cosmos3-Super and Cosmos3-Nano score 67.8 and 63.7 on Physics-IQ Verified. In closed-loop driving, it improves the base policy's distance per at-fault event by 45.5%. In humanoid motion tracking from scratch, it surpasses FPO++ and Gaussian PPO on all seven motions with 33% fewer environment steps.

1. Introduction

Flow models [41, 44] generate images [6, 17], videos [34, 81], driving trajectories [57], and embodied actions [5]. They are often post-trained with RL rewards that are not differentiable. An object detector checks whether an image shows the prompted objects (Fig. 1), and a simulator checks whether a policy completes its task. A differentiable reward gives each sample a gradient to follow [12, 15, 66], whereas a non-differentiable one gives a score but no direction.

Two widely used approaches turn these scores into updates. DiffusionNFT [107] and policy gradients [4, 20, 98] such as Flow-GRPO [43] weight each sample by its own reward. DiffusionART [84] moves all samples of a prompt along one direction, from the mean of their latents toward

their reward-weighted mean. Both approaches pursue the same target (Sec. 3.6), the reward-weighted distribution of a reference model p_{ref} [35],

$$p^{\star}(x \mid c) \propto p_{\mathrm{ref}}(x \mid c)\, \exp\big(r(x,c)/\beta\big), \tag{1}$$

where x is a sample, c its condition (a prompt or a state), r the reward, and β a temperature. A local kernel can instead use each sample's position relative to higher- and lower-reward samples to estimate a direction.

We formulate post-training as representation distribution matching (RDM) [21]: a *group* of samples sharing a prompt or state estimates both the current distribution and its reward-weighted target (Eq. (1)). Matching these distributions in fixed features gives sample directions without reward gradients (Fig. 2). With a linear kernel on latents, this matching gives the direction of DiffusionART, and with rank weights the preference field of DrPO [32] for one-step generators (Sec. 3.6). FlowRDM fits these directions along multi-step sampling trajectories; in control, it forms groups at visited states so that the update needs no action likelihood. The framework separates three choices: **the distribution discrepancy**, **the representation**, and **group formation**. We study these choices across visual generation and embodied control, establishing task-specific recipes within a shared post-training formulation.

The two domains stress different axes. Visual generators readily sample a prompt many times; in the image study, combining the generator's features with its latents improves matching over either representation alone (Sec. 4). In embodied control, the policy acts in closed loop: rollouts from the same start soon reach different states, so a later state holds a single sample, and only the action executed there has an outcome to score, unless the simulator can return to that state. Existing group-based methods therefore compare whole rollouts and trail PPO [75] on robot benchmarks [8, 102]. FlowRDM instead forms a group at each state, where the policy samples several actions. In robotics, a critic with generalized advantage estimation (GAE) [74] scores the executed action, and the others score zero, their expected advantage (Sec. 3.2).

Across the four tasks, FlowRDM improves visual content and closed-loop behavior with the same matching formulation (Fig. 1). On SD3.5-Medium and FLUX.1-dev, it achieves the highest total reward among the compared image post-training methods and also leads on untrained reward models (Sec. 5.1). On Wan2.2 [81], it reaches a VideoReward score [42] of 0.52 against DiffusionART's 0.22 at about 200 training GPU hours (Fig. 3). It also scales to the Cosmos3 world model [58]: post-trained on simulated videos alone, Cosmos3-Super and Cosmos3-Nano reach 67.8 and 63.7 on Physics-IQ Verified [53, 68], against publicly reported base-model scores of 65.5 and 61.8 [63] (Sec. 5.2). In closed-loop driving, it raises distance per at-fault event from 1.89 to 2.75 km for a 2B vision-language-action policy (Tab. 4). In manipulation, it matches or outperforms FPO++ [99] on five tasks, reaching 0.64 success against 0.54 on Square (Tab. 5). In humanoid motion tracking from scratch, it exceeds the 295M-step return of FPO++ and Gaussian PPO within 197M steps on all seven motions (Fig. 5).

2. Related Work

Post-training with black-box rewards. Policy gradients treat the sampler as a Markov decision process [4, 20]; for flow models, group-relative variants [77] make the deterministic sampler stochastic [26, 43, 86, 98]. Likelihood-free methods regress on samples that are reweighted [7, 16, 18, 62, 97], ordered by preference or reward [46, 83, 89], or, in the forward process, split or centered by reward [107, 110]. Tilt Matching [65] and, group by group, DiffusionART [84] move the velocity along a covariance with the reward, the first-order effect of the reweighting. Reward score matching [36] and a path-space view [95] unify many of these methods, and the latter averages the value gradient at a noisy sample over its group with a Gaussian kernel on latents; Sec. 3.6 ties them to one reward-weighted target, and FlowRDM measures its kernel in a representation.

Distribution matching in a representation. The MMD [23, 54] has trained generators [37, 38], and its gradient flow moves samples along the witness gradient [2]; Stein variational policy gradient [45] couples a set of policies through a kernel on their parameters. RDM [21] trains one-step generators with it in the features of frozen pretrained encoders, drifting models [13] and related objectives [10, 60, 103] match generators to data in feature space, and LADD [72] takes its discriminator features, and LPO [104] its reward model, from a diffusion model, as FlowRDM does for images and videos (Sec. 3.4). A reward, where present, is a separate term next to distillation toward a teacher [19, 29, 30, 100] or a regularizer toward the pretrained model [61]. Among feature-space matching methods, DrPO [32] is the most related: for one-step image generators, it ranks the samples of a prompt and adds a kernel field in encoder features, toward high-ranked and away from low-ranked samples, to a drift from the frozen base generator. Its preference field is the witness gradient of Sec. 3.3 with rank weights (Sec. 3.6); Tab. 1 ablates its normalized form within our multi-step training loop. FlowRDM connects these fields to a reward-weighted target, fits their sample directions along multi-step trajectories, and extends the grouping to states of closed-loop policies.

Closed-loop policies. In closed loop, each state of a rollout holds one executed action, which online RL for diffu-

Algorithm 1 One round of FlowRDM

Require: velocity field v_θ, fixed representation ϕ, discrepancy D, reward r, group size K, step size η

$v_{\text{old}} \leftarrow v_\theta$
for each condition c of the round **do** ▷ Sec. 3.2
  sample $x_{1:K}$ with v_{old} given c; cache the trajectories
  score $r_{1:K}$; skip c if all rewards are equal
  $\alpha_k \propto r_k - \bar{r}$; $f_k \leftarrow \phi(x_k)$
  $d_i \leftarrow -\nabla_{f_i} D(\hat{p}_c, \hat{p}_c^\star)$ ▷ Eq. (4); MMD: Eq. (9)
  $g_i \leftarrow J_\phi(x_i)^\top d_i$ or g_c, rescaled ▷ Eqs. (5) and (6)
  targets $v_{\text{old}}(x_{i,t_n}, t_n \mid c) - \eta\, g_i$ for all i, n
end for
fit v_θ to the targets ▷ Eq. (7)

sion and flow policies scores with a critic or the rollout outcome [14, 47, 48, 51, 67, 71, 87, 99, 105], and group-relative updates, also used for driving planners [39, 108], compare rollouts that share only an initial state [80, 109]. DPMD [47] fits a diffusion policy to the target of Eq. (1), with a critic as the reward, by reweighted score matching. MPO [1] compares several samples at one state, and the vine procedure of TRPO [73] executes several actions from a restored simulator state, as FlowRDM does in driving. FlowRDM forms its groups at states (Sec. 3.2) and matches them in a representation, which needs no policy likelihood.

3. FlowRDM

FlowRDM post-trains a flow model in rounds; Fig. 2 illustrates one round. In each round, the model forms groups (Sec. 3.2), the rewards reweight the samples into a target, matching in fixed features gives the directions (Secs. 3.3 and 3.4), and the network update fits shifted velocity targets (Sec. 3.5). Algorithm 1 summarizes one round, and Tab. 6 in Appendix F lists the choices of each task.

3.1. Setup and target

A flow model generates a sample by integrating a learned velocity field $v_\theta(x_t, t \mid c)$ from Gaussian noise at $t = 1$ to data at $t = 0$ [41, 44]. The condition c is the input that the model is given: a text prompt for an image or video generator, joined by the input video in video-to-video generation, and the state for a policy, whose sample x is an action or a sequence of actions. With an N-step sampler on times $1 = t_N > \cdots > t_0 = 0$, an Euler step reads

$$x_{t_{n-1}} = x_{t_n} + (t_{n-1} - t_n)\, v_\theta(x_{t_n}, t_n \mid c), \qquad (2)$$

and $p_\theta(x \mid c)$ denotes the distribution of the sample $x = x_{t_0}$. A reward $r(x, c)$ scores a sample; it can be a learned scorer or the outcome of a simulation.

In a round, the reference in Eq. (1) is the current model, $p_{\text{ref}} = p_{\text{old}}$, with velocity v_{old} and parameters θ_{old}.

3.2. Group formation and selection

Formation. Matching compares the samples of a *group*, K samples $x_1, \ldots, x_K$ of the current model that share a condition c, with rewards $r_k = r(x_k, c)$. A visual generator forms such a group by sampling one condition K times. An embodied policy acts in closed loop and visits each state once, so we form the group at a state of a rollout, where the policy draws K samples. In robotics, only the executed action receives its critic-based generalized advantage estimate (GAE) [74]; the other samples receive zero, using the policy's zero expected advantage as a baseline for their unobserved outcomes. With advantages as rewards, the target of Eq. (1) is the improved policy of MPO [1] and AWR [62].

Selection. A group with equal rewards carries no direction, since reweighting leaves it unchanged (Appendix D). Sampling needs no backward pass and can be much cheaper than training, so group selection spends samples and updates on conditions with unequal rewards (Appendix F), as RL for language models does [96, 101, 106].

3.3. Matching a group to its reward-weighted group

The target $p^\star$ is available only through samples of the current model and their rewards. Since a group comes from the reference p_{old}, self-normalized importance weights turn it into an estimate of the target on the same samples,

$$\hat{p}_c = \frac{1}{K} \sum_k \delta_{x_k}, \qquad \hat{p}_c^\star = \sum_k w_k\, \delta_{x_k}, \qquad (3)$$

with $w_k = \operatorname{softmax}_k(r/\beta)$, and we call $\hat{p}_c^\star$ the *reward-weighted group*. Its weights differ from the uniform ones by the *centered weights* $\alpha_k = w_k - 1/K$, which sum to zero. To first order in $1/\beta$, $\alpha_k \approx (r_k - \bar{r})/(K\beta)$, with $\bar{r}$ the mean reward of the group, and FlowRDM always uses this first-order form. The reward enters only through these weights and the model only through its samples, so the matching needs no reward gradient and no sampler likelihood.

FlowRDM matches the two groups in fixed features $f_k = \phi(x_k)$ (Sec. 3.4), holding the reward-weighted group fixed during the update. For squared maximum mean discrepancy (MMD) D with kernel κ [23], its target-minus-source witness u_c [54] and feature descent direction are

$$\begin{aligned} u_c(f) &= \sum_k \alpha_k\, \kappa(f, f_k), \\ d_i &= \nabla_f u_c(f_i) \propto -\nabla_{f_i} D(\hat{p}_c, \hat{p}_c^\star). \end{aligned} \qquad (4)$$

Thus the centered rewards define a differentiable kernel field without differentiating the reward (Appendix A). Under a Gaussian kernel, each sample moves toward nearby samples with above-average reward and away from nearby samples with below-average reward. Other differentiable discrepancies give analogous descent directions;

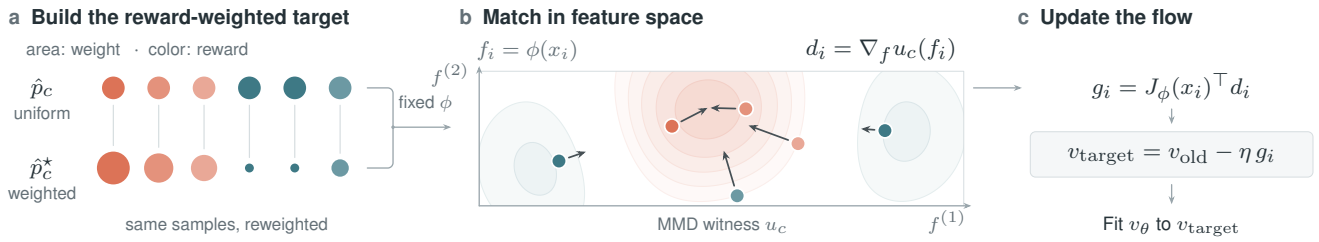

Figure 2. **FlowRDM turns reward scores into matching directions.** (a) Reweighting the same six samples by reward defines a fixed target. Circle area is probability mass; warm and cool colors mark above- and below-mean rewards. (b) In fixed features $f_i = \phi(x_i)$, shading shows the MMD witness u_c and arrows its gradients d_i. (c) Pulling these directions back gives g_i, which shifts the velocity targets that fit the model. The example uses per-sample directions, a Gaussian kernel, positive first-order weights, and one common arrow scale.

we also study a drifting field [13] and use (de)-regularized MMD [11] for video.

3.4. The representation and the matching direction

Matching in a representation, as in RDM [21] and drifting models [13], adds an inductive bias: the representation ϕ, any fixed map, chosen per task (Tab. 6), from a sample (possibly with its condition) to features, decides which samples count as near, and the matching direction is informative when near samples have similar rewards. In visual generation, the curse of dimensionality makes samples nearly equidistant [3], so a kernel on raw samples loses its power [69], whereas in learned features, such as those of a self-supervised encoder [59] or of the generator itself [72], near samples share content. In our image study, generator features improve both linear and Gaussian kernels over their latent-space counterparts (Sec. 4). In embodied control, actions and plans are low-dimensional, yet a prior still helps: the Frenet frame of the route [90] gives a plan's progress and lateral offset (Sec. 5.3).

The Jacobian of a differentiable ϕ carries the feature direction d_i back to the sample as the *matching direction*

$$g_i = J_\phi(x_i)^\top d_i, \tag{5}$$

at the cost of one vector-Jacobian product. Without a derivative of ϕ, the covariance over the group between the witness values $u_k = u_c(f_k)$ and the samples gives the whole group one direction (Appendix A),

$$g_c \propto \sum_k (u_k - \bar{u})\,(x_k - \bar{x}), \tag{6}$$

with $\bar{u}$ and $\bar{x}$ the group means. When rewards are compared only within a group, each group's directions are rescaled to unit root mean square, so the step size alone sets their length and, for the MMD, takes the place of β together with the number of rounds (Appendix C); an advantage, with one scale across states, keeps its length.

3.5. Network update

The matching direction says where each sample should move, and the network update makes the velocity field produce that move. An adjoint pass [9, 15] carries g_i through the sampler (Appendix E); in our 50-update comparison it costs 1.5× per update without improving the score (Tab. 1). We therefore adopt a regression update in the style of Tilt Matching [65] and DiffusionART [84]: at cached points x_{i,t_n} of the sampling trajectory of each sample x_i, the network regresses toward the velocity of the current model shifted against the matching direction,

$$\mathcal{L}(\theta) = \sum_{i,n} t_n \left\| v_\theta(x_{i,t_n}, t_n \mid c_i) - \big(v_{\text{old}}(x_{i,t_n}, t_n \mid c_i) - \eta\, g_i\big) \right\|^2, \tag{7}$$

where c_i is the condition of x_i and η is a step size. Descending-time integration (Eq. (2)) gives the minus sign: holding the cached trajectory fixed, the velocity shift accumulates to an endpoint displacement ηg_i.

3.6. Relation to existing methods

Policy gradients [4, 43] and DiffusionNFT [107] pursue the target of Eq. (1) without letting the samples of a group interact (Appendix B). Policy gradients descend the KL divergence to this target; their group estimate also has a finite-sample algebraic correspondence with a delta-kernel MMD plug-in update (Appendix B). DiffusionNFT regresses on re-noised samples toward a positive distribution that is the same target to first order. DiffusionART lets samples interact only through their latent mean.

Proposition 1 (Mean matching)**.** *With the identity representation on latents and the linear kernel $\kappa(x, y) = x^\top y$, the MMD gives every sample the same direction $d_i = \sum_k \alpha_k x_k$, the group direction of DiffusionART [84], and Eq. (7) with this direction is its update.*

The Gaussian kernel on latents tends to this direction as its bandwidth grows. DrPO [32] weights each sample by wins minus losses in reward-ordered pairs; with them, the witness gradient is its preference field, computed in encoder features for one-step generators (Appendix B).

Configuration	GenEval	GPU h
Representation, kernel, and field (endpoint direction at every sampler state)		
latent, linear kernel (ART direction)	0.909	92
latent, Gaussian kernel	0.918	93
DiT text features + latent, linear kernel	0.931	115
DiT text features + latent, drifting field	0.901	118
DiT text features, Gaussian kernel, no latent	0.905	118
DiT text features + latent, Gaussian kernel (FlowRDM, ours)	0.941	116
DiT text + image features + latent, Gaussian kernel	0.931	127
Group selection		
FlowRDM (ours) with group selection	**0.967**	118
Gradient path		
adjoint pass through the sampler†	0.803	0.42/upd.
forward process (re-noised endpoints)	0.924	141
DiffusionART, released recipe	0.924	97
DiffusionNFT, released recipe	0.965	204

Table 1. **Design-space study on text-to-image generation** with SD3.5-Medium and the GenEval reward, one seed per row. GenEval: mean reward on 2212 evaluation images after 500 updates; GPU h: training compute on 24 H100 GPUs. **Bold**: best. †Image-stream features, 50 updates (Appendix F).

4. Studying the Design Space

On text-to-image generation, where choices can be varied in isolation, we vary the kernel and the representation and check the gradient path and group selection; group formation matters in closed loop and is set in Secs. 5.3 and 5.4.

Setup. We post-train SD3.5-Medium [17] with LoRA [28] and the GenEval reward [22] (rollouts and updates in Appendix F).

Representation. On flattened latents, the Gaussian and linear kernels score 0.918 and 0.909 after 500 updates. We add the prompt-token states of the frozen SD3.5 transformer, read from the sampled image at three noise levels; joint attention makes these features depend on the image. Adding these features raises the linear-kernel score to 0.931 and the Gaussian-kernel score to 0.941. Over the last five evaluations, the Gaussian kernel on these features scores 0.941 to 0.961 and the linear one 0.917 to 0.938. The drifting field in the same representation scores 0.901. The latent entry contributes a tenth of the direction's energy: removing it reduces the Gaussian score from 0.941 to 0.905.

Group selection. With the GenEval reward alone, more and more groups have equal rewards and carry no direction: after 300 updates, 13 of the 48 groups of a round still train. Group selection samples the prompts of unequal groups again in place of fresh prompts, with new noise (Sec. 3.2), and keeps 26 groups training. The run trails the plain one by up to 0.02 for the first 300 updates, matches it from 350 to 450, and ends at 0.967 where the plain run falls back to 0.941: the level of DiffusionNFT at 0.58× its compute.

5. Results Across Domains

We test whether one distribution-matching objective improves generated content and closed-loop actions, with task-specific representations, kernels, and groups (Tab. 6); baselines are re-run from the same base model unless noted.

5.1. Text-to-image generation

We post-train SD3.5-Medium and FLUX.1-dev [6] in the multi-reward setting of DiffusionART [84]: GenEval, PickScore [33], CLIPScore [27], and HPSv2.1 [93] on the GenEval training prompts, with equal weights on fixed-normalized scores. The total reward, the sum of the four normalized scores, is the training reward, evaluated on four images for each of the 553 GenEval evaluation prompts. DiffusionART and DiffusionNFT [107] run the recipes released for each model, on FLUX.1-dev both from the DiffusionART release; training compute is given for 24 H100 GPUs (Appendix F). Here and in Sec. 5.2, every reward is used as a black box, differentiable or not.

SD3.5-Medium. With the recipe of Sec. 4, FlowRDM raises the total reward from the base model's 2.014 to 3.215 after 500 updates (128 GPU hours; Figs. 3 and 6). DiffusionNFT and DiffusionART peak at 3.121 and 3.100 after 200 updates and then decline, to 3.110 and 2.993 at the training compute of FlowRDM. On untrained rewards, FlowRDM also scores highest: ImageReward [94] is 1.55 against 1.39 for DiffusionART and 1.14 for DiffusionNFT, and UnifiedReward [88] is 4.46 against 4.22 and 4.35.

FLUX.1-dev. On FLUX.1-dev (12B), DiffusionART reports its largest-model result. FlowRDM keeps the model, the LoRA, the prompts, the group size, the reward weights, and the regression of DiffusionART's FLUX recipe and changes the direction and the rollouts: a Gaussian-kernel witness in the image-stream features of the FLUX transformer itself, on 10-step stochastic rollouts instead of 6-step deterministic ones (Appendix F). As a regularizer, the loss carries a KL term that holds the mean token of the denoised estimate at the statistics of the base model's samples.

FlowRDM reaches the highest total reward, 3.198 after 500 updates (205 GPU hours; 3.188 without KL), against 3.106 for DiffusionART after 600 updates and 2.995 for DiffusionNFT after 50 updates (Tab. 2 and Fig. 3). From 200 updates on, every evaluation of FlowRDM is above the best of DiffusionART's 1000-update run. FlowRDM also scores highest on ImageReward and UnifiedReward, which are not trained (Tab. 2).

5.2. Video generation

Text-to-video. We post-train Wan2.2-TI2V-5B [81] for text-to-video generation with the VideoReward reward [42]

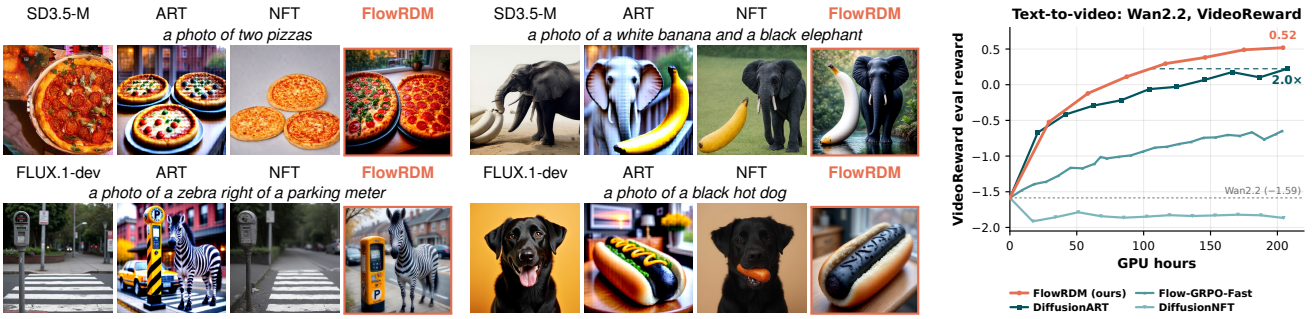

Figure 3. **Visual generation.** Left: four rewards, the same prompt and evaluation sample for every method; top: SD3.5-Medium (base without CFG) after 500 updates (more in Fig. 7); bottom: FLUX.1-dev at the evaluations of Tab. 2. Right: Wan2.2 with VideoReward on 100 held-out prompts; the dashed line marks where FlowRDM first reaches DiffusionART's final reward.

		Trained rewards					Not trained	
Model	GPU h	GenEval	PickScore	CLIPScore	HPSv2.1	Total	ImageReward	UnifiedReward
FLUX.1-dev	–	0.654	23.45	28.29	0.304	2.803	0.95	4.11
+ DiffusionNFT [107]	39	<u>0.797</u>	23.55	29.87	0.296	2.995	1.19	<u>4.43</u>
+ DiffusionART [84]	302	0.756	<u>24.07</u>	<u>32.24</u>	**0.350**	<u>3.106</u>	<u>1.53</u>	4.34
+ FlowRDM (ours)	205	**0.817**	**24.67**	**32.55**	<u>0.348</u>	**3.198**	**1.54**	**4.44**

Table 2. **Four-reward post-training of FLUX.1-dev** on 2212 evaluation images. Total, the training reward, sums the four fixed-normalized trained rewards; the others are raw scores. Each row is the evaluation with the highest Total of its run. **Bold**: best; <u>underline</u>: second best.

and the training prompts of DiffusionART. FlowRDM uses Eq. (5) with a Gaussian kernel on the latents and on the features of the Wan transformer (Appendix F). DiffusionART runs in the same loop with its group direction, so the two differ only in the direction. Flow-GRPO-Fast [43] and DiffusionNFT [107] use the settings that DiffusionART reports for HunyuanVideo. We evaluate the reward on 100 prompts excluded from training and compare the methods at equal training compute.

At the budget of about 200 training GPU hours (204 for FlowRDM and 207 for DiffusionART), FlowRDM raises the reward from -1.59 to 0.52 and DiffusionART to 0.22 (Fig. 3, right). Paired over the 100 prompts, FlowRDM ends 0.29 above DiffusionART (95% bootstrap interval 0.20 to 0.38), and it passes DiffusionART's final reward after 105 GPU hours, about half the budget. Flow-GRPO-Fast reaches -0.65, and DiffusionNFT stays below the base model. FlowRDM leads in visual quality and text alignment (Tab. 7 in Appendix F).

Wan2.2 at its release setting. We also post-train Wan2.2 at the setting of its release, 121 frames at 1280×704 with 50 UniPC steps and guidance 5, on the prompts of Movie Gen Video Bench [64], with HPSv3 [49] and the text alignment of VideoReward as the reward (Appendix F). On 128 held-out prompts, a vision-language model, GPT-6-Astra, that sees the videos of two models in random order names the better one (protocol and instruction in Appendix F).

After 24 updates, the judge prefers FlowRDM to the base model in 78% of the 256 pairs (95% interval 73 to 83; Tab. 3). VideoScore2 [25], not part of the reward, rises from 4.52 to 4.78 in visual quality, from 4.13 to 4.25 in text alignment, and from 4.10 to 4.33 in physical consistency, on scales of 1 to 5. DiffusionART runs in the same loop with the same reward and settings (Appendix F). After the same 24 updates, the judge prefers FlowRDM to DiffusionART in 62% of 128 pairs, one for each prompt (53 to 70).

Preferred	Overall	Visual	Motion	Text
FlowRDM (ours) over Wan2.2	78 (73–83)	87	65	70
FlowRDM (ours) over DiffusionART [84]	62 (53–70)	60	49	64

Table 3. **Blind pairwise preference for Wan2.2 at its release setting, after 24 updates.** Share of pairs (%) in which the vision-language judge, GPT-6-Astra, prefers the video of the first model, overall and for visual quality, motion, and text alignment, on 128 prompts excluded from training; a tie counts as half, and the parentheses give 95% bootstrap intervals. First row: 256 pairs from two seeds with the inference code released with Wan2.2. Second row: 128 pairs, the two directions trained with identical settings (Appendix F).

World model. We post-train Cosmos3 [58], Nano and Super, for video-to-video generation and evaluate on Physics-IQ [53], where the model continues a three-second input video for five seconds and is scored against the real continuation by three motion-mask scores and a pixel-error score. The training set contains no video or prompt of the benchmark: it is 336 simulated rigid-body clips from the physical

interaction scenes released with Cosmos3, and the reward of a sample is the benchmark's score against the simulated continuation of its clip. We compare FlowRDM, Eq. (5) with a Gaussian kernel on the latents and on features of the Cosmos3 transformer, with DiffusionART's group direction (the linear kernel on latents, Proposition 1) under the same settings (Appendix F).

On Cosmos3-Nano, after 32 updates, FlowRDM raises the Physics-IQ Verified score [68], the mean of the four scores, from 55.79 ± 0.91 to 60.84 ± 0.24 over four sampling seeds, and DiffusionART to 58.44 ± 0.93 (paired difference 2.40, 95% bootstrap interval 1.47 to 3.36; Tab. 8). On Cosmos3-Super, eight updates raise the score from 58.89 ± 1.43 to 63.64 ± 0.68. Averaging the weights [70, 92] of three runs that differ in the kernel gives 64.16 ± 0.48.

Best-of-eight evaluation. The Cosmos3 and Odyssey entries in Fig. 1 take one of eight samples per item, chosen with a reward model or by consensus. For the weight-averaged models, we choose each sample by its agreement with the other seven, without ground truth. This gives 67.85 for Cosmos3-Super and 63.69 for Cosmos3-Nano; the best public entry on the leaderboard [63] has 66.10, and the base models 65.47 and 61.83 (64.95 and 62.19 with our selection).

5.3. Closed-loop autonomous driving

Setup. The policy is an Alpamayo-like vision-language-action (VLA) policy [57] with 2B parameters, trained on proprietary driving data. Its trajectory head is a flow-matching expert, which turns noise into a 6.4 s plan; we post-train the expert with the vision-language backbone frozen. The conditions are pre-failure states, shortly before closed-loop failures in simulation [56]. At each state, the policy samples 16 plans, and the vehicle controller executes every plan from that state without replanning. This open-loop outcome estimates each plan's closed-loop outcome, so every plan of a group is scored directly, with no advantage estimate (Sec. 3.2). The reward is 1 for a plan that causes no failure and keeps at least 95% of the recorded progress, and 0 for a failure. Group selection (Sec. 3.2) keeps the 1,788 states with both outcomes (Appendix F).

The representation expresses each plan in the Frenet frame of the route [90], as arc length and signed lateral offset at eight horizons between 0.5 and 6.4 s. The kernel is linear, so the witness function is the projection of these features onto the difference between the mean features of passing and failing plans, and the matching direction is the group form of Eq. (6). Flow-GRPO [43], FPO++ [99], and DiffusionNFT [107] make one pass over the same 25,662 scored plans, with the same backbone and seeds and the base policy as reference (Appendix F).

We evaluate in closed loop on 441 held-out scenarios:

Method	Score	km/AF	Off-road	Collision	Rear-end
Alpamayo 1.5 (10B) [55]	0.794	1.59	28	24	32
Base policy (2B)	0.840	1.89	30	15	23
Flow-GRPO [43]	0.832	1.80	36	11	19
FPO++ [99]	0.836	1.98	34	9	29
DiffusionNFT [107]	0.832	1.75	34	14	27
FlowRDM (ours): Frenet, linear kernel	**0.879**	**2.75**	25	7	25
Changes to FlowRDM (ours)					
Gaussian kernel	0.845	2.21	29	10	25
Rewards as witness	0.850	2.12	26	14	34
Cartesian waypoints	0.854	2.35	27	10	23
Passing plans only	0.857	2.27	28	10	21

Table 4. **Closed-loop driving** on 441 held-out scenarios. Score is the mean scenario score, km/AF the distance per at-fault event; the other columns count scenarios with a road departure, at-fault collision, or rear-end collision by another vehicle. Each row is the run with the higher score of two training seeds; the base policy and Alpamayo 1.5 are the highest-scoring of three evaluations.

the driving score averages route progress over the scenarios, with zero for a scenario with a failure, and km/AF is the distance per at-fault event (Appendix F).

Results. FlowRDM raises the base policy's score from 0.840 to 0.879 and km/AF from 1.89 to 2.75 (Tab. 4). Flow-GRPO, FPO++, and DiffusionNFT stay near the base policy, with at most 0.836 and 1.98 km. In the same evaluation, the public Alpamayo 1.5 (10B) [55] scores 0.794 with a km/AF of 1.59, both below the 2B base policy, and FlowRDM reaches 1.7 times that km/AF.

Ablations. Each change in the lower block of Tab. 4 scores below FlowRDM, and forming the groups at states on recorded trajectories or on failure-free rollouts keeps the score at the base policy's level (Tab. 10 in Appendix F).

5.4. Robot control

Manipulation. We use the fine-tuning suite of FPO++ [99], with Can and Square from RoboMimic [50] and Box Cleanup, Tray Lift, and Threading from DexMimicGen [31]. Each task starts from an image-based flow-matching policy pretrained on demonstrations, which predicts chunks of 16 actions, and we run the released FPO++ training loop with its base policies, critic, and optimizer and replace its policy loss. Groups are formed as in Sec. 3.2: at each state of a rollout, the policy draws $K = 9$ chunks and executes one, which receives its generalized advantage estimate [74]; the other eight receive zero and need no environment steps. The representation is the normalized action chunk and the kernel is linear, so the group direction is $A\,(x_\mathrm{e} - \bar{x}_{-\mathrm{e}})$, with x_e the executed chunk, A its advantage, and $\bar{x}_{-\mathrm{e}}$ the mean of the other eight (Proposition 1): the policy's mean at the state moves along the advantage, as a Gaussian policy's mean does under PPO, with no likelihood or probability ratio.

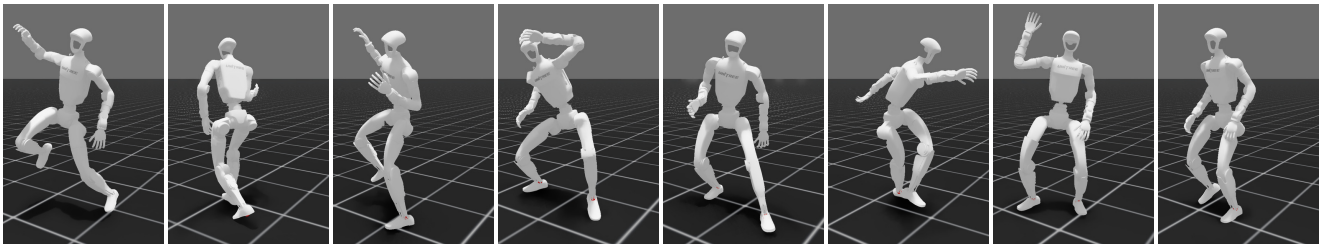

Figure 4. **The Unitree G1 on Dance 2** under a policy trained from scratch with FlowRDM for 197M environment steps: one episode with zero-noise sampling, with frames 1.4 s apart from left to right.

		FPO++ [99]		
Task	Steps	reported	re-run	FlowRDM (ours)
Can	5.0M	0.95	0.939 ± 0.003	$\mathbf{0.986 \pm 0.001}$
Square	8.0M	0.54	0.550 ± 0.002	$\mathbf{0.637 \pm 0.008}$
Box Cleanup	5.0M	1.00	$0.994 \pm 0.005^{\dagger}$	$\mathbf{1.000 \pm 0.000}$
Tray Lift	4.3M	0.95	0.932 ± 0.007	$\mathbf{0.937 \pm 0.009}$
Threading	8.0M	0.89	0.852 ± 0.020	$\mathbf{0.923 \pm 0.012}$

Table 5. **Manipulation fine-tuning** on the suite of FPO++: success rate of the in-training evaluation over the last 40% of the steps in the second column, with the standard error over three seeds. Reported is the mean curve of the FPO++ paper over the same window. **Bold**: higher of re-run FPO++ and FlowRDM. †The evaluations of these runs end at 77 to 87% of the range.

Averaged over the last 40% of each reported curve and three seeds (Tab. 5; Fig. 8 in Appendix F), FlowRDM reaches 0.986 on Can and 0.637 on Square, against 0.95 and 0.54 for the reported FPO++ curve and 0.939 and 0.550 for our FPO++ runs in the same loop, and every seed of FlowRDM is above both. On Threading it reaches 0.923, against 0.89 and 0.852, and on Tray Lift and Box Cleanup the two are level.

Motion tracking. The Unitree G1 humanoid tracks seven motions of LAFAN1 [24] in the BeyondMimic environment [40], six of them as in the FPO++ paper, with one policy per motion trained from scratch (Fig. 4). At each state, the policy draws eight single actions in addition to the executed one, as in manipulation, and each sample gets its own direction, with a kernel that adds a quadratic term per action dimension (Appendix F). FlowRDM trains in fewer environment steps than both baselines: its zero-noise evaluation return after 197M steps is above the 295M-step return of Gaussian PPO [75] in BeyondMimic's released configuration and of released FPO++ on all seven motions (Fig. 5). After 295M steps it averages 36.7 on the six motions of the FPO++ paper, against 33.3 and 33.9 (Tab. 11 in Appendix F).

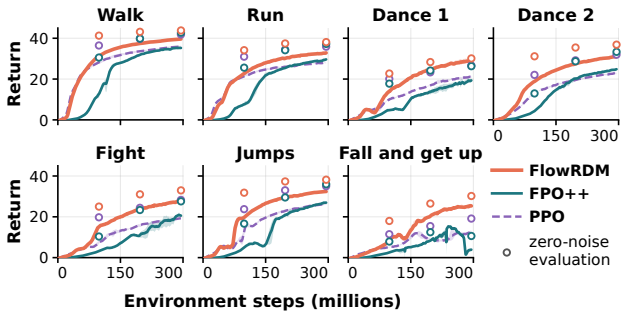

Figure 5. **Humanoid motion tracking from scratch** on the Unitree G1, with three seeds per method. Lines: mean training return, with the range of the seeds shaded. Circles: mean zero-noise evaluation return after about 98M, 197M, and 295M steps. FPO++ and PPO are our runs of the released configurations. On Fall and get up, FPO++ uses the surviving seeds: two after 137M steps, one after 236M.

6. Discussion

The experiments suggest how to set the three choices of FlowRDM for a new task. The *feature space* should expose reward-relevant differences; for high-dimensional images and videos, informative features offer a route to more stable and efficient matching (Sec. 3.4 and Tab. 1). The *distribution discrepancy* should reflect the desired correction: a linear kernel shifts the action mean, while a quadratic term also adjusts its spread (Eq. (10)). In control, *group formation* must compare actions at the same state to turn reward or advantage contrasts into directions. When sampling and scoring are cheaper than training and most groups have uniform rewards, prioritize *group selection*, as in driving.

We hypothesize that this geometric view also explains why FlowRDM transfers to embodied control. Even low-dimensional action spaces can have narrow feasible regions: centimeter-scale trajectory changes may separate success from failure. DiffusionNFT's mirrored velocity regression [107] and GRPO's suppression of low-reward samples [43] do not explicitly use pairwise action geometry to determine a nearby correction. FlowRDM instead couples reward contrasts with the geometry of actions at the same state to construct an explicit corrective displacement.

7. Conclusion

FlowRDM post-trains a flow model by matching each group of samples to its reward-weighted version in fixed features, which gives every sample a direction without reward gradients or sampler likelihoods. DiffusionART's direction and DrPO's preference field are special cases of this matching. With per-task representations, kernels, and groups, the same matching improves image, video, and world-model generators, a closed-loop driving policy, and robot policies that match or exceed FPO++ in manipulation and humanoid tracking from scratch.

Limitations. Representations and kernels are chosen per task. Closed-loop grouping needs a critic or a simulator that can revisit states.

LLM Usage Statement

We used LLM-based coding agents (Claude Code with Claude models, Anthropic; Codex, OpenAI) throughout this work, under the direction of the authors. **Code and experiments.** The agents wrote and modified training and evaluation code, launched and monitored the runs of Secs. 4 and 5, and collected their results. **Text.** The agents drafted and edited all sections of the paper, including the appendix and the derivations, from the authors' outlines and instructions. **Figures and tables.** The agents wrote the plotting scripts and the figure sources and filled the tables from the result files. **Related work.** The agents searched for related papers and summarized them. The authors conceived the method, designed the experiments, and reviewed the code, the results, the references, and the text, and they are responsible for the content of the paper.

References

[1] Abbas Abdolmaleki, Jost Tobias Springenberg, Yuval Tassa, Remi Munos, Nicolas Heess, and Martin Riedmiller. Maximum a posteriori policy optimisation. In *International Conference on Learning Representations*, 2018. 3

[2] Michael Arbel, Anna Korba, Adil Salim, and Arthur Gretton. Maximum mean discrepancy gradient flow. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 2, 14

[3] Kevin Beyer, Jonathan Goldstein, Raghu Ramakrishnan, and Uri Shaft. When Is "Nearest Neighbor" Meaningful? In *International Conference on Database Theory (ICDT)*, 1999. 4

[4] Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training Diffusion Models with Reinforcement Learning. In *International Conference on Learning Representations (ICLR)*, 2024. 1, 2, 4

[5] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. π_0: A Vision-Language-Action Flow Model for General Robot Control. In *Robotics: Science and Systems (RSS)*, 2025. 1

[6] Black Forest Labs. FLUX.1. `https://github.com/black-forest-labs/flux`, 2024. 1, 5

[7] Maksim Bobrin, Maksim Zhdanov, and Dmitry Dylov. Fenchel Tilting: Weighted Correction for Efficient Finetuning of Generative Models. *arXiv preprint arXiv:2609.40030*, 2026. 2

[8] Kang Chen, Zhihao Liu, Tonghe Zhang, Zhen Guo, Si Xu, Hao Lin, Hongzhi Zang, Xiang Li, Quanlu Zhang, Zhaofei Yu, Guoliang Fan, Tiejun Huang, Yu Wang, and Chao Yu. π_{RL}: Online RL fine-tuning for flow-based vision-language-action models. *arXiv preprint arXiv:2510.25889*, 2025. 2

[9] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural Ordinary Differential Equations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. 4

[10] Xingyu Chen, Fei Ma, and Sipei Zhao. DriftAudio: Marginal Drifting for Distributional Post-Training of One-Step Text-to-Audio Generators. *arXiv preprint arXiv:2609.27598*, 2026. 2

[11] Zonghao Chen, Aratrika Mustafi, Pierre Glaser, Anna Korba, Arthur Gretton, and Bharath K. Sriperumbudur. (de)-regularized maximum mean discrepancy gradient flow. *Journal of Machine Learning Research*, 2025. 4, 17

[12] Kevin Clark, Paul Vicol, Kevin Swersky, and David J Fleet. Directly Fine-Tuning Diffusion Models on Differentiable Rewards. In *International Conference on Learning Representations (ICLR)*, 2024. 1

[13] Mingyang Deng, He Li, Tianhong Li, Yilun Du, and Kaiming He. Generative Modeling via Drifting. *arXiv preprint arXiv:2602.04770*, 2026. 2, 4, 15, 16

[14] Shutong Ding, Ke Hu, Shan Zhong, Haoyang Luo, Weinan Zhang, Jingya Wang, Jun Wang, and Ye Shi. GenPO: Generative diffusion models meet on-policy reinforcement learning. In *Advances in Neural Information Processing Systems*, 2025. 3

[15] Carles Domingo-Enrich, Michal Drozdzal, Brian Karrer, and Ricky T. Q. Chen. Adjoint Matching: Fine-tuning Flow and Diffusion Generative Models with Memoryless Stochastic Optimal Control. In *International Conference on Learning Representations (ICLR)*, 2025. 1, 4

[16] Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao, Jipeng Zhang, Kashun Shum, and Tong Zhang. RAFT: Reward rAnked FineTuning for Generative Foundation Model Alignment. *Transactions on Machine Learning Research*, 2023. 2

[17] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling Rectified Flow Transformers for High-Resolution Image Synthesis. In *International Conference on Machine Learning (ICML)*, 2024. 1, 5

[18] Jiajun Fan, Shuaike Shen, Chaoran Cheng, Yuxin Chen, Chumeng Liang, and Ge Liu. Online reward-weighted fine-tuning of flow matching with Wasserstein regularization.

In *International Conference on Learning Representations*, 2025. 2

[19] Linqian Fan, Peiqin Sun, Tiancheng Wen, Shun Lu, and Chengru Song. R_{dm}: Re-conceptualizing distribution matching as a reward for diffusion distillation. In *European Conference on Computer Vision (ECCV)*, 2026. 2

[20] Ying Fan, Olivia Watkins, Yuqing Du, Hao Liu, Moonkyung Ryu, Craig Boutilier, Pieter Abbeel, Mohammad Ghavamzadeh, Kangwook Lee, and Kimin Lee. DPOK: Reinforcement Learning for Fine-tuning Text-to-Image Diffusion Models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 1, 2

[21] Lan Feng, Wuyang Li, Eloi Zablocki, Matthieu Cord, and Alexandre Alahi. Representation Distribution Matching for One-Step Visual Generation. *arXiv preprint arXiv:2607.02375*, 2026. 2, 4

[22] Dhruba Ghosh, Hanna Hajishirzi, and Ludwig Schmidt. GenEval: An Object-Focused Framework for Evaluating Text-to-Image Alignment. In *Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*, 2023. 5

[23] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A Kernel Two-Sample Test. *Journal of Machine Learning Research*, 13 (25):723–773, 2012. 2, 3, 14

[24] Félix G Harvey, Mike Yurick, Derek Nowrouzezahrai, and Christopher Pal. Robust motion in-betweening. *ACM Transactions on Graphics (TOG)*, 39(4):60–1, 2020. 8, 22

[25] Xuan He, Dongfu Jiang, Ping Nie, Minghao Liu, Zhengxuan Jiang, Mingyi Su, Wentao Ma, Junru Lin, Chun Ye, Yi Lu, et al. VideoScore2: Think before you score in generative video evaluation. *arXiv preprint arXiv:2509.22799*, 2025. 6, 19

[26] Xiaoxuan He, Siming Fu, Yuke Zhao, Wanli Li, Jian Yang, Dacheng Yin, Fengyun Rao, and Bo Zhang. TempFlow-GRPO: When timing matters for GRPO in flow models. In *International Conference on Learning Representations (ICLR)*, 2026. 2

[27] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. CLIPScore: A Reference-free Evaluation Metric for Image Captioning. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2021. 5

[28] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations (ICLR)*, 2022. 5

[29] Yushi Huang, Xiangxin Zhou, Ruoyu Wang, Chi Zhang, Jun Zhang, and Tianyu Pang. Reinforcing Few-step Generators via Reward-Tilted Distribution Matching. *arXiv preprint arXiv:2605.26108*, 2026. 2

[30] Dengyang Jiang, Dongyang Liu, Zanyi Wang, Qilong Wu, Liuzhuozheng Li, Hengzhuang Li, Xin Jin, David Liu, Changsheng Lu, Zhen Li, et al. Distribution Matching Distillation Meets Reinforcement Learning. In *European Conference on Computer Vision (ECCV)*, 2026. 2

[31] Zhenyu Jiang, Yuqi Xie, Kevin Lin, Zhenjia Xu, Weikang Wan, Ajay Mandlekar, Linxi Jim Fan, and Yuke Zhu. DexMimicGen: Automated data generation for bimanual dexterous manipulation via imitation learning. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2025. 7

[32] Zhou Jiang, Yandong Wen, and Zhen Liu. Drifting Preference Optimization for One-Step Generative Models. *arXiv preprint arXiv:2606.02521*, 2026. 2, 4, 15, 16

[33] Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. Pick-a-Pic: An Open Dataset of User Preferences for Text-to-Image Generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 5

[34] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, et al. HunyuanVideo: A Systematic Framework For Large Video Generative Models. *arXiv preprint arXiv:2412.03603*, 2024. 1

[35] Tomasz Korbak, Ethan Perez, and Christopher L Buckley. RL with KL penalties is better viewed as Bayesian inference. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, 2022. 2

[36] Jeongjae Lee, Jinho Chang, Jeongsol Kim, and Jong Chul Ye. Reward score matching: Unifying reward-based fine-tuning for flow and diffusion models. *arXiv preprint arXiv:2604.17415*, 2026. 2

[37] Chun-Liang Li, Wei-Cheng Chang, Yu Cheng, Yiming Yang, and Barnabás Póczos. MMD GAN: Towards Deeper Understanding of Moment Matching Network. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 2

[38] Yujia Li, Kevin Swersky, and Richard Zemel. Generative Moment Matching Networks. In *International Conference on Machine Learning (ICML)*, 2015. 2

[39] Yongkang Li, Kaixin Xiong, Xiangyu Guo, Fang Li, Sixu Yan, Gangwei Xu, Lijun Zhou, Long Chen, Haiyang Sun, Bing Wang, Kun Ma, Guang Chen, Hangjun Ye, Wenyu Liu, and Xinggang Wang. ReCogDrive: A reinforced cognitive framework for end-to-end autonomous driving. In *International Conference on Learning Representations (ICLR)*, 2026. 3

[40] Qiayuan Liao, Takara E Truong, Xiaoyu Huang, Yuman Gao, Guy Tevet, Koushil Sreenath, and C Karen Liu. BeyondMimic: From motion tracking to versatile humanoid control via guided diffusion. *Science Robotics*, 11(117): eadx8924, 2026. 8, 21

[41] 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 (ICLR)*, 2023. 1, 3

[42] Jie Liu, Gongye Liu, Jiajun Liang, Ziyang Yuan, Xiaokun Liu, Mingwu Zheng, Xiele Wu, Qiulin Wang, Menghan Xia, Xintao Wang, et al. Improving Video Generation with Human Feedback. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. 2, 5, 17, 18

[43] Jie Liu et al. Flow-GRPO: Training flow matching models

via online rl. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. 1, 2, 4, 6, 7, 8, 14, 18

[44] 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 (ICLR)*, 2023. 1, 3

[45] Yang Liu, Prajit Ramachandran, Qiang Liu, and Jian Peng. Stein variational policy gradient. In *Conference on Uncertainty in Artificial Intelligence*, 2017. 2

[46] Yihong Luo, Tianyang Hu, and Jing Tang. Reinforcing diffusion models by direct group preference optimization. In *International Conference on Learning Representations (ICLR)*, 2026. 2

[47] Haitong Ma, Tianyi Chen, Kai Wang, Na Li, and Bo Dai. Efficient online reinforcement learning for diffusion policy. In *International Conference on Machine Learning*, 2025. 3

[48] Haitong Ma, Chenxiao Gao, Tianyi Chen, Na Li, and Bo Dai. GeMPO: Generalized Measure Matching for Online Diffusion Reinforcement Learning. *arXiv preprint arXiv:2603.10250*, 2026. 3

[49] Yuhang Ma, Yunhao Shui, Xiaoshi Wu, Keqiang Sun, and Hongsheng Li. HPSv3: Towards wide-spectrum human preference score. In *International Conference on Computer Vision (ICCV)*, 2025. 6, 18

[50] Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, and Roberto Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. In *Conference on Robot Learning (CoRL)*, 2021. arXiv:2108.03298. 7

[51] David McAllister, Songwei Ge, Brent Yi, Chung Min Kim, Ethan Weber, Hongsuk Choi, Haiwen Feng, and Angjoo Kanazawa. Flow matching policy gradients. In *International Conference on Learning Representations (ICLR)*, 2026. 3, 21

[52] Mayank Mittal, Pascal Roth, James Tigue, Antoine Richard, Octi Zhang, et al. Isaac Lab: A GPU-accelerated simulation framework for multi-modal robot learning. *arXiv preprint arXiv:2511.04831*, 2025. 21

[53] Saman Motamed, Laura Culp, Kevin Swersky, Priyank Jaini, and Robert Geirhos. Do generative video models understand physical principles? In *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 948–958, 2026. 2, 6

[54] Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, and Bernhard Schölkopf. Kernel Mean Embedding of Distributions: A Review and Beyond. *Foundations and Trends in Machine Learning*, 10(1–2):1–141, 2017. 2, 3, 14

[55] NVIDIA. Alpamayo 1.5. `https://huggingface.co/nvidia/Alpamayo-1.5-10B`, 2026. 7, 21

[56] NVIDIA, Yulong Cao, Riccardo de Lutio, Sanja Fidler, Guillermo Garcia Cobo, Zan Gojcic, Maximilian Igl, Boris Ivanovic, Peter Karkus, Janick Martinez Esturo, Marco Pavone, Aaron Smith, Ellie Tanimura, Michal Tyszkiewicz, Michael Watson, Qi Wu, and Le Zhang. AlpaSim: A modular, lightweight, and data-driven research simulator for autonomous driving. `https://github.com/NVlabs/alpasim`, 2025. 7

[57] NVIDIA, Yan Wang, Wenjie Luo, Junjie Bai, Yulong Cao, Tong Che, Ke Chen, Yuxiao Chen, Jenna Diamond, Yifan Ding, Wenhao Ding, Liang Feng, Greg Heinrich, Jack Huang, Peter Karkus, Boyi Li, Pinyi Li, Tsung-Yi Lin, Dongran Liu, Ming-Yu Liu, Langechuan Liu, Zhijian Liu, Jason Lu, Yunxiang Mao, Pavlo Molchanov, Lindsey Pavao, Zhenghao Peng, Mike Ranzinger, Ed Schmerling, Shida Shen, Yunfei Shi, Sarah Tariq, Ran Tian, Tilman Wekel, Xinshuo Weng, Tianjun Xiao, Eric Yang, Xiaodong Yang, Yurong You, Xiaohui Zeng, Wenyuan Zhang, Boris Ivanovic, and Marco Pavone. Alpamayo-r1: Bridging reasoning and action prediction for generalizable autonomous driving in the long tail, 2025. 1, 7, 21

[58] NVIDIA et al. Cosmos 3: Omnimodal world models for physical AI. *arXiv preprint arXiv:2606.02800*, 2026. 2, 6, 19

[59] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024. 4

[60] Woobin Park, Jaeyeon Kim, Yilun Du, Sitan Chen, and Bohyung Han. Are We Overthinking One-Step Generation? *arXiv preprint arXiv:2610.20331*, 2026. 2

[61] Romeo Passaro, Jeongjae Lee, Jong Chul Ye, Michael M. Bronstein, Yoshua Bengio, Alexander Tong, and Danyal Rehman. Distribution-Based Regularization for Reward Fine-Tuning Diffusion and Flow Models. *arXiv preprint arXiv:2610.24897*, 2026. 2

[62] Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-Weighted Regression: Simple and Scalable Off-Policy Reinforcement Learning. *arXiv preprint arXiv:1910.00177*, 2019. 2, 3

[63] Physics-IQ Verified. Leaderboard. `https://github.com/google-deepmind/physics-IQ-benchmark`, 2026. Video-to-video entries, snapshot of October 20, 2026. 2, 7, 23

[64] Adam Polyak, Amit Zohar, Andrew Brown, Andros Tjandra, Animesh Sinha, Ann Lee, Apoorv Vyas, Bowen Shi, Chih-Yao Ma, Ching-Yao Chuang, et al. Movie Gen: A cast of media foundation models. *arXiv preprint arXiv:2410.13720*, 2024. 6, 18

[65] Peter Potaptchik, Cheuk-Kit Lee, and Michael S. Albergo. Tilt matching for scalable sampling and fine-tuning. In *International Conference on Machine Learning (ICML)*, 2026. 2, 4

[66] Mihir Prabhudesai, Anirudh Goyal, Deepak Pathak, and Katerina Fragkiadaki. Aligning Text-to-Image Diffusion Models with Reward Backpropagation. *arXiv preprint arXiv:2310.03739*, 2023. 1

[67] Michael Psenka, Alejandro Escontrela, Pieter Abbeel, and Yi Ma. Learning a diffusion model policy from rewards via Q-score matching. In *International Conference on Machine Learning*, 2024. 3

[68] Tim Rädsch, Yuki M. Asano, Hilde Kuehne, Stefan Bauer, Priyank Jaini, Robert Geirhos, and Carsten T. Lüth. Physics-IQ Verified. *arXiv preprint arXiv:2606.18943*, 2026. 2, 7

[69] Aaditya Ramdas, Sashank J. Reddi, Barnabás Póczos, Aarti Singh, and Larry Wasserman. On the Decreasing Power of Kernel and Distance Based Nonparametric Hypothesis Tests in High Dimensions. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2015. 4

[70] Alexandre Ramé, Johan Ferret, Nino Vieillard, Robert Dadashi, Léonard Hussenot, Pierre-Louis Cedoz, Pier Giuseppe Sessa, Sertan Girgin, Arthur Douillard, and Olivier Bachem. WARP: On the benefits of weight averaged rewarded policies. *arXiv preprint arXiv:2406.16768*, 2024. 7

[71] Allen Z. Ren, Justin Lidard, Lars L. Ankile, Anthony Simeonov, Pulkit Agrawal, Anirudha Majumdar, Benjamin Burchfiel, Hongkai Dai, and Max Simchowitz. Diffusion policy policy optimization. In *International Conference on Learning Representations (ICLR)*, 2025. arXiv:2409.00588. 3, 21

[72] Axel Sauer, Frederic Boesel, Tim Dockhorn, Andreas Blattmann, Patrick Esser, and Robin Rombach. Fast high-resolution image synthesis with latent adversarial diffusion distillation. In *SIGGRAPH Asia 2024 Conference Papers*, 2024. 2, 4

[73] John Schulman, Sergey Levine, Philipp Moritz, Michael I. Jordan, and Pieter Abbeel. Trust Region Policy Optimization. In *International Conference on Machine Learning (ICML)*, 2015. 3

[74] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. In *International Conference on Learning Representations (ICLR)*, 2016. 2, 3, 7

[75] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal Policy Optimization Algorithms. *arXiv preprint arXiv:1707.06347*, 2017. 2, 8, 21, 22

[76] Clemens Schwarke, Mayank Mittal, Nikita Rudin, David Hoeller, and Marco Hutter. RSL-RL: A learning library for robotics research. *arXiv preprint arXiv:2509.10771*, 2025. 22

[77] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, et al. DeepSeekMath: Pushing the Limits of Mathematical Reasoning in Open Language Models. *arXiv preprint arXiv:2402.03300*, 2024. 2

[78] Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum Likelihood Training of Score-Based Diffusion Models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 15

[79] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei A. Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *International Conference on Machine Learning*, 2020. 19

[80] Shuhan Tan, Kairan Dou, Yue Zhao, and Philipp Krähenbühl. Interactive post-training for vision-language-action models. *arXiv preprint arXiv:2505.17016*, 2025. 3

[81] Team Wan, Ang Wang, Baole Ai, Bin Wen, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025. 1, 2, 5

[82] Manan Tomar, Lior Shani, Yonathan Efroni, and Mohammad Ghavamzadeh. Mirror Descent Policy Optimization. In *International Conference on Learning Representations (ICLR)*, 2022. 15

[83] Bram Wallace, Meihua Dang, Rafael Rafailov, Linqi Zhou, Aaron Lou, Senthil Purushwalkam, Stefano Ermon, Caiming Xiong, Shafiq Joty, and Nikhil Naik. Diffusion model alignment using direct preference optimization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2

[84] Zhenglin Wan, Jiayu Chen, Zhengfeiyang Zhang, Pengfei Zhou, Wangbo Zhao, Fangyuan Kong, Pengfei Wan, Kun Gai, and Yang You. DiffusionART: Efficient, Stable and Scalable Reinforcement of Diffusion Models. *arXiv preprint arXiv:2610.12572*, 2026. 1, 2, 4, 5, 6, 14, 15, 18, 19, 20

[85] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *International Conference on Learning Representations*, 2021. 19

[86] Jing Wang, Jiajun Liang, Jie Liu, Henglin Liu, Gongye Liu, Jun Zheng, Wanyuan Pang, Ao Ma, Zhenyu Xie, Xintao Wang, Meng Wang, Pengfei Wan, and Xiaodan Liang. GRPO-Guard: Mitigating implicit over-optimization in flow matching via regulated clipping. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026. 2

[87] Siting Wang, Xiaofeng Wang, Zheng Zhu, Minnan Pei, Xinyu Cui, Cheng Deng, Jian Zhao, Guan Huang, Haifeng Zhang, and Jun Wang. π-StepNFT: Wider Space Needs Finer Steps in Online RL for Flow-based VLAs. *arXiv preprint arXiv:2603.02083*, 2026. 3

[88] Yibin Wang, Yuhang Zang, Hao Li, Cheng Jin, and Jiaqi Wang. Unified reward model for multimodal understanding and generation. *arXiv preprint arXiv:2503.05236*, 2025. 5

[89] Yiyang Wang, Xi Chen, Xiaogang Xu, Yu Liu, and Hengshuang Zhao. GDRO: Group-level reward post-training suitable for diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026. 2

[90] Moritz Werling, Julius Ziegler, Sören Kammel, and Sebastian Thrun. Optimal trajectory generation for dynamic street scenarios in a frenet frame. In *2010 IEEE international conference on robotics and automation*, pages 987–993. IEEE, 2010. 4, 7

[91] Daan Wierstra, Tom Schaul, Tobias Glasmachers, Yi Sun, Jan Peters, and Jürgen Schmidhuber. Natural Evolution Strategies. *Journal of Machine Learning Research*, 15(27): 949–980, 2014. 22

[92] Mitchell Wortsman, Gabriel Ilharco, Samir Yitzhak Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S. Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International Conference on Machine Learning*, 2022. 7

[93] Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li. Human Preference Score v2: A Solid Benchmark for Evaluating Human Preferences of Text-to-Image Synthesis. *arXiv preprint arXiv:2306.09341*, 2023. 5

[94] Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. ImageReward: Learning and evaluating human preferences for text-to-image generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 5

[95] Yixian Xu, Yuanrui Zhang, Shengjie Luo, Kaiyuan Gao, Liwei Wang, and Di He. Designing Reinforcement Learning for Diffusion Models: A Unified Path-Space View. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2026. 2

[96] Yixuan Even Xu, Yash Savani, Fei Fang, and J. Zico Kolter. Not All Rollouts are Useful: Down-Sampling Rollouts in LLM Reinforcement Learning. *Transactions on Machine Learning Research*, 2026. 3

[97] Shuchen Xue, Chongjian Ge, Shilong Zhang, Yichen Li, and Zhi-Ming Ma. Advantage weighted matching: Aligning RL with pretraining in diffusion models. In *International Conference on Machine Learning*, 2026. 2

[98] Zeyue Xue, Jie Wu, Yu Gao, Fangyuan Kong, Lingting Zhu, Mengzhao Chen, Zhiheng Liu, Wei Liu, Qiushan Guo, Weilin Huang, et al. DanceGRPO: Unleashing GRPO on Visual Generation. *arXiv preprint arXiv:2505.07818*, 2025. 1, 2

[99] Brent Yi, Hongsuk Choi, Himanshu Gaurav Singh, Xiaoyu Huang, Takara E. Truong, Carmelo Sferrazza, Yi Ma, Rocky Duan, Pieter Abbeel, Guanya Shi, Karen Liu, and Angjoo Kanazawa. Flow policy gradients for robot control. In *Conference on Robot Learning (CoRL)*, 2026. 2, 3, 7, 8, 20, 21, 22, 23

[100] Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T. Freeman, and Taesung Park. One-step Diffusion with Distribution Matching Distillation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2

[101] Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, et al. DAPO: An Open-Source LLM Reinforcement Learning System at Scale. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. 3

[102] Hongzhi Zang, Mingjie Wei, Si Xu, Yongji Wu, Zhen Guo, Yuanqing Wang, Hao Lin, Peihong Wang, Liangzhi Shi, Yuqing Xie, Zhexuan Xu, Zhihao Liu, Kang Chen, Wenhao Tang, Quanlu Zhang, Weinan Zhang, Chao Yu, and Yu Wang. RLinf-VLA: A unified and efficient framework for reinforcement learning of vision-language-action models. In *Robotics: Science and Systems (RSS)*, 2026. 2

[103] Jinhua Zhang, Yisong Lin, Wei Long, and Shuhang Gu. Beyond Token-Level Cross-Entropy: Fréchet Distributional Post-Training for Autoregressive Image Generation. *arXiv preprint arXiv:2608.00562*, 2026. 2

[104] Tao Zhang, Cheng Da, Kun Ding, Huan Yang, Kun Jin, Yan Li, Tingting Gao, Di Zhang, Shiming Xiang, and Chunhong Pan. Diffusion model as a noise-aware latent reward model for step-level preference optimization. In *Advances in Neural Information Processing Systems*, 2025. 2

[105] Tonghe Zhang, Chao Yu, Sichang Su, and Yu Wang. ReinFlow: Fine-tuning Flow Matching Policy with Online Reinforcement Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. 3

[106] Haizhong Zheng, Yang Zhou, Brian R. Bartoldson, Bhavya Kailkhura, Fan Lai, Jiawei Zhao, and Beidi Chen. Act Only When It Pays: Efficient Reinforcement Learning for LLM Reasoning via Selective Rollouts. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. 3

[107] Kaiwen Zheng et al. DiffusionNFT: Online diffusion reinforcement with forward process. In *International Conference on Learning Representations (ICLR)*, 2026. 1, 2, 4, 5, 6, 7, 8, 14, 15, 18

[108] Zewei Zhou, Tianhui Cai, Seth Z. Zhao, Yun Zhang, Zhiyu Huang, Bolei Zhou, and Jiaqi Ma. AutoVLA: A vision-language-action model for end-to-end autonomous driving with adaptive reasoning and reinforcement fine-tuning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. 3

[109] Fangqi Zhu, Zhengyang Yan, Zicong Hong, Quanxin Shou, Xiao Ma, and Song Guo. WMPO: World model-based policy optimization for vision-language-action models. In *International Conference on Learning Representations (ICLR)*, 2026. 3

[110] Yuanzhi Zhu, Xi Wang, Stéphane Lathuilière, and Vicky Kalogeiton. Diffusion reinforcement learning via centered reward distillation. *arXiv preprint arXiv:2603.14128*, 2026. 2

[111] Yuxin Zuo, Kaiyan Zhang, Li Sheng, Shang Qu, Ganqu Cui, Xuekai Zhu, Haozhan Li, Yuchen Zhang, Xinwei Long, Ermo Hua, Biqing Qi, Youbang Sun, et al. TTRL: Test-time reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. 19

A. The MMD and Its Witness Function

Between a group and its reward-weighted group, the squared maximum mean discrepancy (MMD) [23] with a kernel κ on features is the quadratic form $\sum_{j,k} \alpha_j \alpha_k \kappa(f_j, f_k)$ in the centered weights, which is large when the reward separates the group in the representation. Its witness function [54],

$$u_c(f) = \sum_k \alpha_k \, \kappa(f, f_k), \tag{8}$$

is the reward of the group, centered and smoothed by the kernel. Unlike the reward, it can be differentiated, and the direction of Eq. (4) is its gradient at the sample's own features, a step of the MMD gradient flow toward the reward-weighted group [2],

$$d_i = \nabla_f u_c(f_i) = \sum_k \alpha_k \, \nabla_1 \kappa(f_i, f_k), \tag{9}$$

where ∇_1 is the gradient in the first argument. For the Gaussian kernel $\kappa(f, f') = \exp\big(-\|f - f'\|^2/(2h^2)\big)$ with bandwidth h, $d_i = h^{-2} \sum_k \alpha_k \, \kappa(f_i, f_k) \, (f_k - f_i)$: each sample moves toward nearby samples with above-average reward and away from nearby samples with below-average reward. The group form of Eq. (6) is, to leading order in the spread of the group, the direction of Eq. (5) at the group mean, multiplied by the sample covariance of the group.

B. Proofs

Proof of Proposition 1. For $\kappa(x, y) = x^\top y$, $\nabla_1 \kappa(x_i, x_k) = x_k$, so Eq. (9) gives $d_i = \sum_k \alpha_k \, x_k$ for every i. With the identity feature map, $J_\phi = I$ and $g_i = d_i$. With the first-order centered weights, $\sum_k \alpha_k x_k = \frac{1}{K\beta} \sum_k (r_k - \bar{r}) \, x_k$ is the group's empirical covariance between reward and latent, divided by β. DiffusionART derives this direction from its Lemma 2, by which the reweighting moves the endpoint mean along the reward-endpoint covariance to first order [84]. It normalizes the direction to unit root mean square, $G(c)$, and regresses the velocity toward $v_{\text{old}} + G(c)$ at every cached state with the loss weight t. Its velocity $u_t = x_0 - \epsilon$ points from noise to data and v_θ of Eq. (2) from data to noise, so in our convention this target is v_{old} minus a multiple of $G(c)$, and the update is Eq. (7).

Proposition 2 (Policy gradient)**.** *The KL-regularized objective $\mathcal{J}(\theta) = \mathbb{E}_{p_\theta}[r] - \beta \, \mathrm{KL}(p_\theta \,\|\, p_{\text{ref}})$ equals $-\beta \, \mathrm{KL}(p_\theta \,\|\, p^\star)$ up to a constant, so ascending it descends the KL divergence to the target of Eq.* (1)*. At $\theta = \theta_{\text{old}}$ with $p_{\text{ref}} = p_{\text{old}}$, the score-function estimate of its gradient on a group, $\sum_i (r_i - \bar{r}) \, \nabla_\theta \log p_\theta(x_i)$, is up to a positive factor the plug-in score-function estimate of $-\nabla_\theta \mathrm{MMD}^2_\kappa(p_\theta, \hat{p}^\star_c)$ under the delta kernel $\kappa(x, y) = \mathbf{1}[x = y]$. Dividing the advantages $r_i - \bar{r}$ by the group's standard deviation, as GRPO does, sets the temperature of each condition to $\beta_c \propto \mathrm{std}_c(r)$.*

KL-regularized objective. Let $p^\star = p_{\text{ref}} \, e^{r/\beta}/Z$ with $Z = \mathbb{E}_{p_{\text{ref}}}[e^{r/\beta}]$. Then $\log p^\star = \log p_{\text{ref}} + r/\beta - \log Z$, and

$$\mathrm{KL}(p_\theta \,\|\, p^\star) = \mathrm{KL}(p_\theta \,\|\, p_{\text{ref}}) - \tfrac{1}{\beta} \mathbb{E}_{p_\theta}[r] + \log Z.$$

Multiplying by $-\beta$ gives $\mathcal{J}(\theta) = \mathbb{E}_{p_\theta}[r] - \beta \, \mathrm{KL}(p_\theta \,\|\, p_{\text{ref}}) = -\beta \, \mathrm{KL}(p_\theta \,\|\, p^\star) + \beta \log Z$.

Proof of Proposition 2. Since $\nabla_\theta \mathrm{KL}(p_\theta \| p_{\text{ref}}) = \mathbb{E}_{p_\theta}[\log(p_\theta/p_{\text{ref}}) \, \nabla_\theta \log p_\theta]$ and the log-ratio vanishes at $\theta = \theta_{\text{old}}$ when $p_{\text{ref}} = p_{\text{old}}$, the score-function gradient of $\mathcal{J}$ there is $\mathbb{E}_{p_\theta}[(r - b) \nabla_\theta \log p_\theta]$ for any baseline b; the group mean gives $\frac{1}{K} \sum_i (r_i - \bar{r}) \nabla_\theta \log p_\theta(x_i)$. For a fixed target $\hat{p}^\star$, let $u(x) = \mathbb{E}_{y \sim \hat{p}^\star} \kappa(x, y) - \mathbb{E}_{y \sim p_\theta} \kappa(x, y)$. The log-derivative trick applied to the two terms of the MMD that depend on θ gives

$$\nabla_\theta \mathrm{MMD}^2_\kappa(p_\theta, \hat{p}^\star) = -2 \, \mathbb{E}_{x \sim p_\theta} \big[u(x) \, \nabla_\theta \log p_\theta(x) \big].$$

With the group in place of p_θ, the witness values are $u(x_i) = \sum_k \alpha_k \, \kappa(x_i, x_k)$, and the estimate is $-\frac{2}{K} \sum_i u(x_i) \nabla_\theta \log p_\theta(x_i)$. Under the delta kernel $\kappa(x, y) = \mathbf{1}[x = y]$, $u(x_i) = \alpha_i$, since the samples of a flow model are distinct with probability one, and with the first-order centered weights $\alpha_i = (r_i - \bar{r})/(K\beta)$ the descent direction is $\frac{2}{K^2 \beta} \sum_i (r_i - \bar{r}) \nabla_\theta \log p_\theta(x_i)$, a positive multiple of the policy gradient. Replacing β by $\beta_c = \mathrm{std}_c(r)$ gives the standardized advantages of GRPO. It also holds when x_i denotes the sampled trajectory and $p_\theta(x_i)$ its likelihood, which is the form Flow-GRPO estimates [43].

Proposition 3 (Forward-process regression)**.** *With $\hat{A}_k = (r_k - \bar{r})/\tau$ for a normalizer $\tau > 0$, the positive distribution of DiffusionNFT [107] is $\pi^+ \propto \big(1 + \mathrm{clip}(\hat{A}, -1, 1)\big) \, p_{\text{old}}$, which is the target of Eq.* (1) *with $p_{\text{ref}} = p_{\text{old}}$ and $\beta = \tau$ to first order in $1/\tau$. Its loss moves the velocity from v_{old} toward that of π^+ by flow-matching regression on re-noised samples in latent space.*

Proof of Proposition 3. DiffusionNFT defines $o = \frac{1}{2} + \frac{1}{2} \mathrm{clip}(\hat{A}, -1, 1) \in [0, 1]$ with $\hat{A} = (r - \bar{r})/\tau$ and $\pi^+ = o \, p_{\text{old}}/\mathbb{E}_{p_{\text{old}}}[o]$ [107], so $\pi^+ \propto (1 + \mathrm{clip}(\hat{A}, -1, 1)) \, p_{\text{old}}$. For $\tau \geq \sup_x |r(x, c) - \bar{r}|$ the clip is inactive, and π^+ is then the first-order expansion of the exponential reweighting $e^{\hat{A}} p_{\text{old}}/\mathbb{E}_{p_{\text{old}}}[e^{\hat{A}}]$. Since $e^{\hat{A}} \propto e^{r/\tau}$, this reweighting is the target of Eq. (1) with $p_{\text{ref}} = p_{\text{old}}$ and $\beta = \tau$. The DiffusionNFT loss is flow-matching regression of the implicit positive velocity onto samples weighted by o, that is, onto

samples of π^+, plus the mirrored term for π^-. Its minimizer is $v_{\text{old}} + \frac{2}{\beta_{\text{NFT}}}\Delta$ with $\Delta = \alpha_{\text{NFT}}(x_t)(v^+ - v_{\text{old}})$ (Theorems 3.1 and 3.2 of [107]), where v^+ is the velocity of π^+, and $\alpha_{\text{NFT}}(x_t)$ and β_{NFT} are the coefficient and the guidance hyperparameter of DiffusionNFT. Flow-matching regression onto samples of a distribution is the forward-process objective whose minimizer is that distribution's velocity; with likelihood weighting it bounds the forward KL of the corresponding diffusion SDE [78], and DiffusionNFT applies an adaptive weighting to the same regression.

Preference field of DrPO. DrPO [32] draws M pairs from the samples of a prompt, orders each pair by reward, and moves the features of a one-step generator along the gradient of $E(z) = \gamma \sum_{j=1}^{M} \left[\kappa(z, a_j^+) - \kappa(z, a_j^-)\right]$, where a_j^+ and a_j^- are the features of the preferred and of the other sample of pair j. Collecting the terms of each sample gives $E(z) = \gamma \sum_k s_k \, \kappa(z, f_k)$, where s_k is the number of pairs that sample k wins minus the number that it loses, so ∇E is Eq. (9) with $\alpha_k \propto s_k$, and $\sum_k s_k = 0$ as for the centered weights. With pairs drawn uniformly at random, the expectation of s_k is proportional to the number of samples with a lower reward minus the number with a higher reward, a centered rank; for a reward with two values, it is proportional to $r_k - \bar{r}$, the centered weight of FlowRDM to first order. In its experiments, DrPO normalizes the two kernel sums as drifting models do [13], the form of the drifting-field row of Tab. 1, and adds a drift toward samples of the frozen base generator.

C. Centered Weights

With the current model as the reference, the target of each round is the outcome of a mirror-descent step on the expected reward, with the KL divergence to the current model as the proximity term [82, 84]. We use $\alpha_k = (r_k - \bar{r})/((K-1)\, s_c)$, where s_c is the mean absolute deviation of the group's rewards from their mean, and $\alpha = 0$ when $s_c = 0$; this is the first-order form at $\beta_c = (K-1)\, s_c/K$. The witness gradient is linear in α, and the matching directions of each group are rescaled to unit root mean square (Appendix F), so the update depends only on the direction of α. The softmax weights give $\alpha(\beta) = \text{softmax}(r/\beta) - 1/K$, whose direction tends to that of $r - \bar{r}$ as $\beta \to \infty$, so the update of FlowRDM is the limit of the rescaled update under the softmax weights, for which $\hat{p}_c^\star$ is a distribution for every β. For the same reason, β_c does not change the update: the step size η sets how far a round shifts the velocity targets, and with the number of rounds it takes the place of the temperature.

D. Per-Condition Discrepancy

Let p be the model at condition c, and write $\rho = p^\star/p = e^{r/\beta}/Z$ with $Z = \mathbb{E}_p[e^{r/\beta}]$, $\mu_r = \mathbb{E}_p[r]$, and $\sigma^2 = \text{Var}_p(r)$. Then $\log Z = \mu_r/\beta + \sigma^2/(2\beta^2) + O(\beta^{-3})$ and $\mathbb{E}_{p^\star}[r] = \mu_r + \sigma^2/\beta + O(\beta^{-2})$, so $\text{KL}(p^\star \| p) = \mathbb{E}_{p^\star}[r]/\beta - \log Z = \sigma^2/(2\beta^2) + O(\beta^{-3})$.

For the MMD, let ψ be the canonical feature map of κ composed with the representation, which is the representation ϕ of Sec. 3.3 itself for the linear kernel. With $\mu_p = \mathbb{E}_p[\psi]$ and $\mu_{p^\star} = \mathbb{E}_{p^\star}[\psi]$, $\text{MMD}^2(p, p^\star) = \|\mu_{p^\star} - \mu_p\|^2$ in the norm of the reproducing kernel Hilbert space of κ. Then $\mu_{p^\star} - \mu_p = \mathbb{E}_p[(\rho - 1)\, \psi]$ with $\rho - 1 = (r - \mu_r)/\beta + O(\beta^{-2})$, so $\mu_{p^\star} - \mu_p = \text{Cov}_p(r, \psi)/\beta + O(\beta^{-2})$ and $\text{MMD}^2 = \|\text{Cov}_p(r, \psi)\|^2/\beta^2 + O(\beta^{-3})$. By the Cauchy-Schwarz inequality, $\|\text{Cov}_p(r, \psi)\| = \|\mathbb{E}_p[(r - \mu_r)(\psi - \mathbb{E}\psi)]\| \le \sigma\, (\mathbb{E}\|\psi - \mathbb{E}\psi\|^2)^{1/2}$. When the rewards at c are equal, $\rho = 1$ and both discrepancies are zero; in a group with equal rewards, the centered weights, and with them the directions, are zero.

E. Adjoint Variant

The adjoint variant of Sec. 4 carries the matching direction g_i to each sampler state with an adjoint pass. For the step of Eq. (2), the adjoint states satisfy

$$\lambda_{i,n} = \left(I + (t_{n-1} - t_n)\, \partial_x v_\theta\right)^\top \lambda_{i,n-1}, \qquad \lambda_{i,0} = g_i,$$

so that $\lambda_{i,n} = (\partial x_{i,t_0}/\partial x_{i,t_n})^\top g_i$. The variant regresses $v_\theta(x_{i,t_n}, t_n)$ toward $v_{\text{old}}(x_{i,t_n}, t_n) - \eta_n \, \bar{\lambda}_{i,n-1}$, with the rescaled state $\bar{\lambda}_{i,n} = \lambda_{i,n}\, \text{RMS}(g_i)/\text{RMS}(\lambda_{i,n})$ and step sizes η_n that absorb the time weight t_n. At $\theta = \theta_{\text{old}}$, the prediction equals v_{old}, so the gradient of this regression is $2 \sum_{i,n} \eta_n \, (\partial_\theta v_\theta(x_{i,t_n}, t_n))^\top \bar{\lambda}_{i,n-1}$. Backpropagating $F(\theta) = \sum_i \langle g_i, \Phi_\theta(z_i) \rangle$, where $\Phi_\theta(z_i)$ is the sample reached from the noise z_i, through the sampler of Eq. (2) gives $\partial_\theta F = \sum_{i,n} (t_{n-1} - t_n)\, (\partial_\theta v_\theta(x_{i,t_n}, t_n))^\top \lambda_{i,n-1}$. Here $t_{n-1} - t_n = -|t_{n-1} - t_n|$, and $\bar{\lambda}_{i,n-1}$ is $\lambda_{i,n-1}$ times the positive factor $\text{RMS}(\bar{\lambda}_{i,n-1})/\text{RMS}(\lambda_{i,n-1})$. A descent step on this regression and an ascent step on F, which moves each sample along g_i, are therefore both positive combinations of the same terms $-(\partial_\theta v_\theta(x_{i,t_n}, t_n))^\top \lambda_{i,n-1}$, with the weights $2\eta_n \, \text{RMS}(\bar{\lambda}_{i,n-1})/\text{RMS}(\lambda_{i,n-1})$ for the regression and $|t_{n-1} - t_n|$ for backpropagation. For a stochastic step with the injected noise held fixed, v_θ enters the step mean through a negative multiple of the identity, so the same argument holds with $|t_{n-1} - t_n|$ replaced by the magnitude of that multiple and with the adjoint recursion using the Jacobian of the stochastic step.

	Text-to-image	Video generation	Closed-loop driving	Robot control
Condition c	prompt	prompt (and video)	state	state
Group formation	K samples per prompt	K samples per condition	K executed plans per state	K samples, one executed
Reward r	reward models	reward model or benchmark score	pass or fail of the plan	advantage estimate or 0
Representation ϕ	generator features	latents, generator features	Frenet coordinates	action (chunk)
MMD kernel κ	Gaussian	Gaussian	linear	linear or degree two
Direction	per sample (Eq. (5))	per sample (Eq. (5))	per group (Eq. (6))	per group or per sample
Group selection	unequal groups reused	none	both outcomes in the group	none
Update variant	regression	regression	regression	regression (+ JVP)

Table 6. **The terms of FlowRDM in each task.** Video generation covers text-to-video and video-to-video. Section 5 gives the models, rewards, and protocols. Humanoid training adds the JVP trajectory correction of Appendix F.5.

F. Implementation Details

F.1. Text-to-image generation

Training. SD3.5-Medium and FLUX.1-dev are post-trained with LoRA of rank 32 at learning rate 3×10^{-4}. On SD3.5-Medium, each rollout draws 24 samples for each of 48 prompts at 512×512 with a 10-step stochastic sampler (noise level 0.7) and no classifier-free guidance, and is followed by two updates. The centered weights are those of Appendix C.

Representation and direction. On SD3.5-Medium, the representation has four entries: the hidden states of the prompt tokens in the text stream of the transformer at blocks 6, 8, and 12, read on the sample re-noised to the noise levels 0.65, 0.5, and 0.8, respectively, and the flattened latent, weighted to a tenth of the energy of the direction. The image-stream entries of Tab. 1 are the hidden states of blocks 10 and 14 at noise level 0.65, pooled to 16×16 tokens. On FLUX.1-dev, the representation is the image-stream hidden states of blocks 8 and 12 at noise level 0.5, without the latent. Every entry has a Gaussian kernel whose bandwidth is 0.5 times the median pairwise feature distance within the groups of the first rollout and is fixed afterwards. The matching direction of Eq. (5) is computed at the final sample with one vector-Jacobian product through the frozen transformer and scaled to unit root mean square over each group.

Update. The matching direction takes the place of the group direction in the regression of DiffusionART, at all 10 cached states of all samples on SD3.5-Medium and at the first six states of half of the samples of each group on FLUX.1-dev. As in DiffusionART, the loss of each state is weighted by its noise level t_n (Eq. (7)), and the trust region gives a state zero weight when the KL divergence per dimension between the one-step transitions of the current and the old model exceeds $\delta = 10^{-5}$ and the regression would move the velocity further from v_{old}. On FLUX.1-dev, the squared error of each trained state adds the KL term $\frac{c}{64\, t_n^2}\,(\mu - m)^\top S^{-1}(\mu - m)$, where μ is the mean token of the denoised estimate $x_{t_n} - t_n v_\theta$ (the mean over all tokens of each of the 64 latent entries of a token), m and S are the mean and the covariance of the mean tokens of 1,536 samples of the base model, and c is the largest variance of their part that differs between the four positions inside a token.

Baselines. DiffusionART and DiffusionNFT run their released recipes; DiffusionNFT makes one update per rollout and evaluates an exponential moving average of its policy. On FLUX.1-dev, DiffusionART runs its released FLUX recipe (6-step deterministic rollouts, LoRA rank 32, 48 prompts with 24 samples each per rollout, two updates per rollout) for 1000 updates, and DiffusionNFT runs the FLUX recipe released with DiffusionART for its 500 updates at one update per rollout.

Evaluation and compute. Every evaluation draws 512×512 images with a deterministic sampler, with 40 steps and no classifier-free guidance on SD3.5-Medium and with 28 steps at guidance scale 3.5 on FLUX.1-dev. At the last evaluation of each FLUX.1-dev run, FlowRDM is at 3.187 (550 updates, 225 GPU hours), DiffusionART at 3.063 (1000 updates, 532 GPU hours), and DiffusionNFT at 2.874 (500 updates, 428 GPU hours). The training compute counts sampling, scoring, and updates on 24 H100 GPUs and leaves out the evaluations; for FlowRDM on FLUX.1-dev, it is the number of updates times the measured cost of one update on 24 GPUs, 0.409 GPU hours.

Design-space study. In Tab. 1, the ART direction is the Gaussian field in the limit of an infinitely wide kernel, run in the training loop of the other rows. The drifting field of drifting models [13], which DrPO [32] adopts for one-step generators, replaces the two kernel sums of the witness gradient by affinity-normalized means over the above-mean and the below-mean samples of the group. On top of the transformer features, the latent entry adds 0.04 to 0.10 at equal updates, and the image-stream entries give no further gain (0.931). Over the last five checkpoints of one training

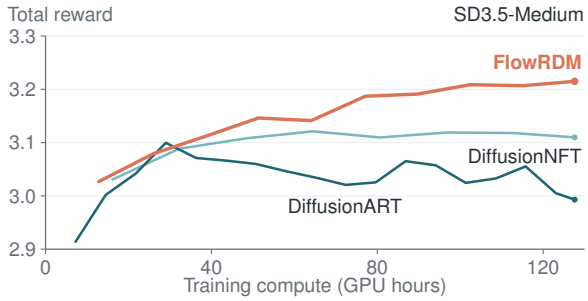

Figure 6. **SD3.5-Medium post-training at equal compute.** Total reward over the training budget of FlowRDM, evaluated on four images for each of the 553 GenEval evaluation prompts. The total sums the four fixed-normalized rewards of Sec. 5.1.

run per configuration (300 to 500 updates), the Gaussian and the linear kernel on the DiT text features and the latent average 0.953 and 0.931, the two kernels on the latent alone 0.912 and 0.895, and the released recipes of DiffusionART and DiffusionNFT 0.895 and 0.958. The adjoint row uses the image-stream features for 50 updates, at 1.5× the cost per update of its no-adjoint twin (0.813 at the same update). Regressing on re-noised endpoints instead of sampler states, the training distribution of DiffusionNFT, reaches 0.924 after 500 updates at 1.2× the cost. Per update, FlowRDM costs 1.2× the released DiffusionART recipe (0.232 against 0.194 GPU hours) and 0.57× DiffusionNFT (0.408); the extra over ART is the truncated passes of the frozen transformer that compute the features of the re-noised endpoints, and their vector-Jacobian products.

Group selection. Prompts whose groups had unequal rewards are sampled again with new noise in the next rollout, with preference for groups whose rewards split evenly around their mean; the compute per update is unchanged. On SD3.5-Medium with the GenEval reward (Tab. 1), from 100 updates on, 20 of the 48 prompts of a round are on average carried over, and 83% of their groups have unequal rewards again, against 37% of the groups of fresh prompts. On FLUX.1-dev with the GenEval reward alone, and with the direction, the rollouts, and the evaluation of Sec. 5.1 without the KL term, 21 of the 48 prompts of a rollout are on average carried over, and 79% of their groups have unequal rewards again, against 24% of the groups of fresh prompts. From 300 to 650 updates, every evaluation of this run lies between 0.966 and 0.986, against 0.955 to 0.970 for the run without selection, which draws 48 fresh prompts in every rollout; after 650 updates, the two runs are at 0.981 and 0.964. The released recipe of DiffusionART peaks at 0.968 with this reward and ends at 0.964 after 1000 updates. With the four rewards of Sec. 5.1, no group has equal rewards, so group selection keeps every group.

F.2. Video generation

Text-to-video. Following DiffusionART, the reward is VideoReward [42] with weights 0.2, 0.1, and 0.7 on its normalized scores of visual quality, motion quality, and text alignment, and the training prompts are those of DiffusionART. Wan2.2-TI2V-5B is fine-tuned in full with AdamW at learning rate 10^{-5} with weight decay 10^{-4}, and gradient clipping at 1. Each round draws 32 samples for each of 32 prompts, 81 frames at 480×832, with a 16-step sampler (shift 5) and no classifier-free guidance. The centered weights are those of Appendix C and, like the witness, use all 32 samples of a group; eight samples of each group are trained, and the two updates of a round take disjoint halves of the prompts. The representation has two families: the latent, a spatially pooled latent, and the temporal mean of a pooled latent; and the hidden states of block 15 of the 30 transformer blocks at noise levels 0.5 and 0.9, as all tokens and as pooled tokens. Every entry has a Gaussian kernel whose bandwidth is 0.5 times the median pairwise distance of the samples of the round. The witness weights are $(\mathbf{K}/K + \lambda I)^{-1}\alpha$ [11] in place of α, where $\mathbf{K}$ is the kernel matrix of the group and λ is three times the mean eigenvalue of $\mathbf{K}/K$. The direction of the feature family is pulled back through the generator, and the directions of the two families are each scaled to unit root mean square over the group and averaged. A second term moves each sample away from the samples of the other prompts in the pooled latent, with a weight of 0.5 times the relative decrease of the median distance between prompts since the first round, so it is zero unless the prompts move closer. The velocity target takes the direction at the 11 highest-noise steps of the sampler.

DiffusionART uses the same loop, with the linear kernel on latents; Flow-GRPO-Fast and DiffusionNFT use the configurations of the DiffusionART paper for HunyuanVideo, with 8 samples for each of 32 prompts on 32 GPUs. Evaluation uses 100 prompts excluded from training and a 30-step sampler (shift 5) with the same initial noise at every evaluation, every 20 updates (every 40 for Flow-GRPO-Fast), and the intervals are 95% bootstrap intervals over the 100 prompts. The training GPU hours count sampling, decoding, scoring, and updates, and leave out the evaluations and the loading of checkpoints when a run resumes; an update costs 1.45 GPU hours for FlowRDM, 1.03 for DiffusionART, 0.23 for Flow-GRPO-Fast, and 0.85 for DiffusionNFT. At equal updates, after the 140 updates of FlowRDM, DiffusionART is at 0.07, which is 0.45 below FlowRDM (0.38 to 0.53).

Wan2.2 at its release setting. The model runs at the configuration of the Wan2.2-TI2V-5B release: 121 frames at 1280×704 and 24 frames per second, 50 steps of UniPC with shift 5, and classifier-free guidance 5 with the release's



Figure 7. **Four-reward post-training of SD3.5-Medium** on six GenEval evaluation prompts (Fig. 3 shows two), with the same evaluation sample for every method: the base model (no CFG) and the three post-trained models after 500 updates.

Method	VQ	MQ	TA	Reward
Wan2.2	−1.05	−0.98	−1.82	−1.59
+ Flow-GRPO-Fast [43]	−0.19	−0.49	−0.80	−0.65
+ DiffusionNFT [107]	−1.00	0.22	−2.41	−1.87
+ DiffusionART [84]	2.42	**1.65**	−0.61	0.22
+ FlowRDM (ours)	**3.16**	1.45	**−0.37**	**0.52**

Table 7. **Text-to-video post-training of Wan2.2 with VideoReward, after 200 training GPU hours.** VQ, MQ, and TA are the normalized VideoReward scores of visual quality, motion quality, and text alignment on the 100 evaluation prompts, and Reward is $0.2\,\mathrm{VQ} + 0.1\,\mathrm{MQ} + 0.7\,\mathrm{TA}$, the training reward. The budget holds 900 updates of Flow-GRPO-Fast, 240 of DiffusionNFT, 200 of DiffusionART, and 140 of FlowRDM. The rows are the last evaluations of the curves of Fig. 3, one training run per method. **Bold**: best.

negative prompt. The rollouts of training use this sampler, so the reward scores the videos that the released pipeline produces. The prompts are the 1,003 prompts of Movie Gen Video Bench [64], extended by the prompt extension of the Wan2.2 release with Qwen2.5-14B-Instruct; 128 of them, stratified over the benchmark's concept and motion labels, are excluded from training and used for evaluation, and the other 875 are the training prompts. The reward has three scores: HPSv3 [49] on eight frames of the video with the content-free prompt "A high-quality image", as the mean over the frames; HPSv3 on the same frames with the video's prompt, as the mean over the three highest-scoring frames; and the text-alignment score of VideoReward [42]. Each score is centered and divided by its standard deviation among base-model samples of the same prompt, and the three enter with weights 0.5, 0.5, and 1. Training is full fine-tuning with the optimizer settings of the text-to-video runs on 128 GPUs. Each round draws 16 samples for each of 64 prompts, trains eight samples of each group, and makes four updates; the velocity target takes the direction at seven of the 32 highest-noise steps of the sampler, drawn for each sample. The direction is that of the world-model runs below, with the hidden states of block 15 of the 30 transformer blocks as generator features, and the loss also carries a penalty on the deviation of the velocity from that of the base model, with an adaptive weight. In the comparison of Tab. 3, FlowRDM and DiffusionART, the linear kernel on latents, are trained in this loop with identical settings; with either direction, 24 updates take about 185 GPU hours, evaluations included.

Pairwise judge. The judge of Tab. 3 is the vision-language model GPT-6-Astra, queried once per pair. For a pair it receives the prompt and one contact sheet per video, 12 frames evenly spaced over the video in time order, in three rows of four frames of 640 pixels width, labeled only A and B in an order drawn at random for each pair, and it answers A, B, or tie for visual quality, motion, text alignment, and the overall preference of a viewer. Its instruction is:

> You are comparing two AI-generated videos made for the same text prompt. Image 1 is a contact sheet of Video A and image 2 is a contact sheet of Video B. Each sheet shows 12 frames of its 5-second video in time order, left to right and top to bottom (first frame top-left, last frame bottom-right). Text prompt given to the video model: "[prompt]". Compare the two

videos and decide, for each criterion, which one is better ("A" or "B"), or "tie" only if you genuinely cannot separate them: *visual_quality*: image quality within the frames: sharpness and fine detail, lighting and colour, and the absence of artifacts (smearing, garbled textures, deformed faces / hands / limbs / objects, duplicated or melting parts). *motion*: what the frame sequence shows over time: do subjects and objects keep their identity and shape, does the motion make physical sense, do things appear or vanish without cause? A video in which nothing moves although the prompt describes motion is worse on this criterion than one with plausible motion. *alignment*: how faithfully the video shows what the prompt asks for (subjects, attributes, action, setting, style). *overall*: which video a careful viewer would prefer, all things considered. Look closely at faces, hands, text and the moving parts before deciding. Answer with the JSON object only.

The preference is the share of pairs won, with a tie as half, and the intervals are 95% bootstrap intervals over the pairs. For the first row of Tab. 3, the base model and the post-trained model generate with the inference code released with Wan2.2 from seeds 0 and 1, the same seed for both videos of a pair; the second row uses the evaluation of the training loop, with the same sampler and one initial noise per prompt for both models. VideoScore2 [25] scores the 256 videos of each model of the first row on scales of 1 to 5; paired over the videos, its gains are 0.26 in visual quality (0.19 to 0.33), 0.12 in text alignment (0.04 to 0.20), and 0.22 in physical consistency (0.13 to 0.31). In the comparison of the second row, the reward on the evaluation prompts rises by 1.62 for FlowRDM and by 1.07 for DiffusionART after 24 updates, in the standardized units of the reward.

World model. The training set has 336 clips from 173 simulations of the physical interaction scenes released with Cosmos3 [58], with bowling, dominoes, billiards, falling objects, ramps, and collapsing towers, each from up to three fixed cameras. Cosmos3 was trained on these scenes; post-training adds the physical reward. A clip has 117 frames at 24 frames per second, the first 21 as the input, and a prompt in the format of the benchmark's prompts, written by GPT-6-Astra from frames of the clip, with two of the benchmark's prompts as examples of the format. Clips are kept when their continuation moves in 0.5% to 18% of the frame, and 24 clips show a scene at rest; both follow the statistics of the benchmark's continuations. The reward of a sample is the benchmark's score against the simulated continuation of its clip; a simulation has one take, so the reward takes the median of the benchmark's physical variances as its tolerance. Both models are fine-tuned in full with AdamW at learning rate 10^{-5} on 128 GPUs, with 64 inputs and four updates per round, Cosmos3-Nano (15B parameters) for 32 updates and Cosmos3-Super (64B) for 8. Each round draws 16 samples for each input with a 16-step sampler at guidance 3 and trains half of the samples of each group, with a penalty on the deviation of the velocity from that of the base model, whose weight is adapted to hold the deviation at 0.12. The latent kernel has three entries, the latent, a spatially pooled latent, and the temporal mean of a pooled latent, and the generator features are the hidden states of the middle transformer block (block 18 of 36 for Nano, block 32 of 64 for Super) at noise levels 0.5 and 0.9, as all tokens and as tokens pooled to $3 \times 3 \times 4$.

Model	ST	S	WS	MSE	Verified
Cosmos3-Super	53.5	71.0	61.2	49.9	58.89 ± 1.43
+ DiffusionART [84]	63.3	72.5	63.6	54.3	63.42 ± 0.93
+ FlowRDM (ours)	**63.6**	72.2	63.9	54.8	63.64 ± 0.68
+ FlowRDM (ours), three runs averaged	**63.6**	**72.7**	**64.7**	**55.7**	$\mathbf{64.16 \pm 0.48}$
Cosmos3-Nano	52.6	67.4	57.3	45.9	55.79 ± 0.91
+ DiffusionART [84]	57.0	69.5	59.3	48.0	58.44 ± 0.93
+ FlowRDM (ours)	**62.7**	70.7	62.8	47.1	60.84 ± 0.24
+ FlowRDM (ours), three runs averaged	60.7	**71.0**	**62.9**	**48.8**	$\mathbf{60.86 \pm 0.57}$

Table 8. **Post-training of Cosmos3 on simulated clips, evaluated on Physics-IQ (Verified, video-to-video).** Eight updates on Cosmos3-Super and 32 on Cosmos3-Nano; the training set contains no video or prompt of the benchmark. ST, S, WS, and MSE are the benchmark's spatiotemporal, spatial, and weighted-spatial mask scores and its pixel-error score (higher is better), and Verified is their mean; all are means over four sampling seeds, the Verified score with its standard deviation. The two post-trained rows of a model differ only in the direction, and its last row is the uniform average of the weights of three runs: these two and one with the Gaussian kernel on latents alone. **Bold**: best for each model.

Evaluation uses the inference code released with Cosmos3 and the release recipe of the benchmark's Cosmos3 entries, with 1280×720, 189 frames, 50 steps of UniPC, guidance 3, shift 4.75, and the upsampled prompts of those entries, and the official evaluator. Single-sample scores are the mean and standard deviation over four runs with different sampling seeds, as the benchmark recommends (Tab. 8). Paired over the 198 items, with each item's score averaged over the four runs, FlowRDM is 5.05 above the base model on Cosmos3-Nano (95% bootstrap interval 3.79 to 6.37) and 2.40 above DiffusionART (1.47 to 3.36); on Cosmos3-Super, FlowRDM is 4.74 above the base model (3.50 to 6.03), and the average of the three runs is 5.27 above the base model (4.06 to 6.51) and 0.53 above FlowRDM alone (0.05 to 1.04). For best-of-eight sampling (Sec. 5.2), the eight samples of an item come from the weight-averaged model with the inference code released with Cosmos3, and the selected sample is the one that agrees most with the other seven under the benchmark's score.

Test-time reinforcement learning. We also post-train Cosmos3-Nano on the prompts and input videos of the 198 test items [111] (cf. test-time adaptation [79, 85]), with the recipe above and a second regularizer, a flow-matching term

Model	ST	S	WS	MSE	Verified
Cosmos3-Nano	51.3	70.8	60.8	46.1	57.24
+ DiffusionART [84]	57.0	70.9	61.5	46.5	58.99
+ FlowRDM (ours), latents only	55.2	71.3	62.4	49.2	59.53
+ FlowRDM (ours)	**62.0**	**72.6**	**64.5**	**50.1**	**62.30**

Table 9. **Test-time post-training of Cosmos3-Nano on Physics-IQ (Verified, video-to-video), six updates each, mean of two training runs.** Columns as in Tab. 8. The post-trained rows differ only in the direction. **Bold**: best.

on four samples of the base model for each input. The reward uses no ground-truth continuation: it is the agreement of a sample's motion masks with those four samples under the benchmark's three mask scores. These runs are sampled with the Cosmos3 pipeline of the diffusers library at one sampling seed, where the base model scores 57.24. After six updates, averaged over two training runs, FlowRDM reaches 62.30 and DiffusionART 58.99, and FlowRDM is highest on all four components (Tab. 9). Paired over the 198 items, FlowRDM is 5.10 and 5.01 above the base model in the two training runs (95% bootstrap intervals 3.34 to 6.83 and 3.26 to 6.75) and 4.09 and 2.53 above DiffusionART (2.64 to 5.57 and 1.14 to 3.99).

F.3. Closed-loop driving

Policy and states. The flow-matching expert has 28 layers of width 512 (199M parameters) and outputs 64 steps of acceleration and curvature at 10 Hz; the vision-language backbone is frozen. The 1,788 states come from closed-loop failures of the base policy and of other policies of the same family in 920 training scenarios, which are disjoint from the 441 evaluation scenarios. At each state, the policy samples 16 plans from one conditioning with a 10-step stochastic sampler (noise level 0.7), and the vehicle controller executes each plan against the map and the logged agents.

Reward. The reward is 1 for a plan that causes no collision, stays on the road and inside the corridor around the recorded route, and keeps at least 95% of the recorded progress, and 0 for a collision, a road departure, a corridor exit, or a stall. Safe plans that keep less than 95% of the recorded progress in a state where another safe plan keeps it, and plans whose replay is invalid, are left out; in the 242 states where no plan keeps 95% of the recorded progress, every safe plan passes. Of the 25,662 scored plans, 8,289 pass and 17,373 fail.

Training. The Frenet features are the arc length relative to the start of the plan, scaled by 0.776, and the signed lateral offset along the recorded route at 0.5, 1, 1.5, 2, 3, 4, 5, and 6.4 s; the scale is fixed on base-policy samples. With the linear kernel, Eq. (9) gives every plan of a group the same witness gradient, and the group form of Eq. (6) carries it to the actions without a derivative of the Frenet map. The group direction has unit root mean square and enters the velocity target with unit gain for every plan of the group, at one of the plan's ten cached sampler states. The loss is weighted by the noise level, v_{old} is the base policy, and the trust region is that of the text-to-image runs, with $\delta = 10^{-5}$. Training is one pass over the scored plans, 101 updates at batch size 256 with AdamW, learning rate 4×10^{-4}, a cosine schedule with 4% warm-up, weight decay 0.1, and gradient clipping at 1.

Ablations. The Gaussian-kernel ablation uses a bandwidth of 3.96 m, 0.6 times the median pairwise feature distance of base-policy samples, and the Cartesian ablation uses the ego-frame waypoint coordinates at the same horizons. The ablation with passing plans only trains the expert with the flow-matching loss on the 8,289 passing plans, with the same 101 updates. The ablation with the rewards as witness uses the binary reward of each plan in place of its witness value in Eq. (6) and needs no representation.

Baselines. All methods train for one pass over the same scored plans with the same seeds, and the baselines use the optimizer settings above except for the learning rate of FPO++. DiffusionNFT uses its loss with $\beta_{\mathrm{NFT}} = 1$ and the binary reward as the optimality label, fresh noise on each plan, no loss weight and no trust region, and the base policy as v_{old}. Flow-GRPO uses its clipped policy-gradient loss on one transition of each plan's stored 10-step chain per pass, with the Gaussian transition of the sampler (noise level 0.7), log-probabilities averaged over the action dimensions, and the old log-probability from the base policy. Its settings are the released ones: clip range 10^{-4}, KL weight 0.04, and advantages equal to the group-centred reward divided by the standard deviation of all rewards, clipped at 5. FPO++ [99] uses the configuration of its authors for fine-tuning with sparse rewards: a per-sample ratio, a clip range of 0.01, the symmetric trust region of PPO in place of the asymmetric one, the advantage of Flow-GRPO, and the learning rate of its authors, 10^{-5}. Its ratio is the exponential of the difference between the flow-matching losses of the base policy and of the current policy on one fresh noise level and noise sample per plan per pass, with the Huber loss ($\delta = 1$) averaged over the action dimensions and summed over the 64 steps.

Evaluation. Evaluation replans at 2 Hz with a 10-step deterministic sampler at initial-noise temperature 0.5, drawing six candidate plans and executing the one closest to the route. Each of the 441 held-out scenarios is rolled out

Groups and training target	Score	km/AF	Off-road	Collision	Rear-end
FlowRDM (ours): pre-failure states	**0.879**	**2.75**	25	7	25
States on recorded trajectories	0.839	2.00	30	12	43
States on rollouts without a failure	0.841	1.91	33	12	32
Recorded trajectory as target	0.577	0.53	69	44	70

Table 10. **Group formation in closed-loop driving**, in the protocol and with the columns of Tab. 4; the run with the higher score of two training seeds. Rows two and three form 1,788 groups at states on replays of the recorded trajectories and on rollouts of the base policy without a failure; the last row trains with the flow-matching loss toward the recorded trajectory at the pre-failure states.

once. A scenario scores 0 when the policy causes a collision, leaves the road, or drifts more than 4 m from the recorded route, and otherwise scores its route progress divided by 0.8 and capped at 1; the driving score is the mean over scenarios. The distance per at-fault event (km/AF) is the total distance over all scenarios divided by the number of scenarios with an at-fault collision or a road departure. Alpamayo 1.5 is evaluated from its public checkpoint with 10B parameters [55, 57], with the same settings on the same scenarios.

Group formation. Table 10 changes where the groups are formed and what the expert is trained toward, with two seeds per row. Two rows form the same number of groups with both outcomes, 1,788, at other states: at fixed times along replays of the recorded trajectories, and every 5 s along rollouts of the base policy that end without a failure. The gain depends on where the groups are formed: at these states, 54% and 57% of the scored plans pass, against 32% at the pre-failure states, and FlowRDM scores 0.839 and 0.841, the level of the base policy. The last row keeps the pre-failure states and trains the expert with the flow-matching loss toward the recorded trajectory in place of the policy's passing plans. It scores 0.577: its route progress is 0.67, against 0.90 for the base policy, and it is rear-ended in 70 of the 441 scenarios.

F.4. Robot manipulation

Training. The runs use the released fine-tuning loop of FPO++ [99] with its base policies, critic, discount, gradient clipping, and optimizer (AdamW, learning rate 10^{-5} for the policy and 10^{-4} for the critic, 10 epochs of 8 minibatches per iteration), and the policy loss of FlowRDM takes the place of the clipped surrogate. An iteration collects 1,600 steps on 30 environments (48,000 environment steps), which gives 3,000 chunks of 16 actions. At the state of each executed chunk, eight further chunks are drawn with the 10-step sampler of the policy, and their sampling trajectories are cached. The reward of the executed chunk is the generalized advantage estimate of its first step ($\lambda = 0.99$), standardized over the iteration and clipped at ± 3.

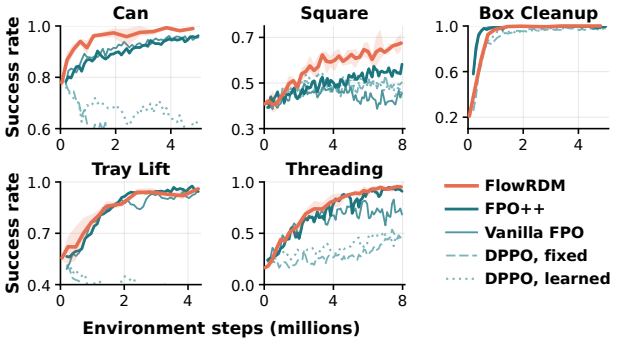

Figure 8. **Manipulation fine-tuning** on the FPO++ suite: success rate of the in-training evaluation against environment steps; FlowRDM: mean and range of three seeds; FPO++, vanilla FPO, and DPPO with fixed and learned noise: mean curves from the FPO++ paper [99].

Update. The group direction $g = A\,(x_{\mathrm{e}} - \bar{x}_{-\mathrm{e}})$ of Sec. 5.4 is computed in normalized action units and is not rescaled per group, so one step size holds for all states. The velocity at one cached point per chunk and minibatch is regressed toward the velocity of the current policy shifted against $\eta\, g$, as in Eq. (7) without the noise-level weight and the trust region, and the loss is 0.2 times the squared error divided by the mean squared velocity change that is asked for. Groups whose executed chunk contains the end of an episode and a held-out tenth of the groups have no policy term, and the first iteration trains only the critic. The per-task hyperparameters are the scale of the sampler's initial noise (0.30 on Can, 1.50 on Square, 0.73 on Box Cleanup, 0.92 on Tray Lift, and 0.66 on Threading) and the step η (0.35 on Square and 0.6 on the other tasks).

Baselines and evaluation. FPO++ weights a probability ratio taken from flow-matching losses with the advantage estimate of the same critic in the clipped surrogate of PPO [75]; we run it with three seeds and its released settings in the same loop. The curves of FPO++, of vanilla FPO [51], and of two DPPO variants are the ones reported in the FPO++ paper, means over three runs, read from the vector figure of its arXiv source. The DPPO variants are implementations of DPPO [71] for flow policies by the FPO++ authors, with a fixed and with a learned exploration noise. The evaluation is the in-training evaluation of the released code, with the sampler started from zero noise, every 5 or 10 iterations. In Tab. 5, Reported is the mean curve of the FPO++ paper averaged over the window at the evaluation steps of FlowRDM, to two decimals; Fig. 8 overlays FlowRDM on the reported curves.

F.5. Humanoid motion tracking

Task. The task is `Tracking-Flat-G1-v0` of BeyondMimic [40] in Isaac Lab [52], as the released Isaac Lab ex-

periments of FPO++ [99] use it. The Unitree G1 tracks one reference motion on flat ground with position targets for its 29 joints at 50 Hz, in episodes of 10 s that start at a phase of the motion drawn by the sampler of the task, which favors phases at which episodes have failed. The seven motions are those of the released code, clips of the LAFAN1 dataset [24] retargeted to the G1 and between 131 and 261 s long: walking, running, two dances, fighting, jumping, and falling and getting up; the FPO++ paper uses the first six. The policy observes 160 values, the reference joint state and proprioception with noise, and the critic observes 286 values without noise. The reward sums six tracking terms and three penalties and is at most 50 per episode. An episode ends early when the height of the torso, a hand, or a foot deviates from the reference by more than 0.25 m or the orientation of the torso deviates beyond a threshold, and the robot is pushed at random every 1 to 3 s.

Direction. For a visited state with executed action x_e and reward A, the rollout policy draws eight actions $x_1, \dots, x_8$ from new noise and caches their sampling trajectories, so each state has nine actions, as in manipulation. For sample k, the group is the executed action and the other seven samples, so that $K = 8$ in what follows, with the perturbation of the rollout added to the samples. Let $\bar{x}_{-k}$ be the mean of the seven samples, and m_k and σ_k^2 the mean and the variance of the K points, per dimension, with the variance normalized by K. The direction of sample k is, per dimension,

$$\begin{aligned} g_k &= A\,(x_\mathrm{e} - \bar{x}_{-k}) + \tfrac{1}{2}\,A\,q_k\,(x_k - m_k), \\ q_k &= \frac{K}{K-1}\Big(\frac{(x_\mathrm{e} - m_k)^2}{\sigma_k^2} - 1\Big). \end{aligned} \tag{10}$$

This is Eq. (5) at x_k for the kernel $\kappa(x, y) = \sum_j \big[\tilde{x}_j \tilde{y}_j + \tilde{x}_j^2\,\tilde{y}_j^2/(4\sigma_{k,j}^2)\big]$ on the centered actions $\tilde{x} = x - m_k$, with the weight A on the executed action and $-A/(K-1)$ on each of the other seven samples, so that sample k is left out of its own group. The first term moves the mean of the policy at the state, and the second changes its spread: for a positive reward, the samples move away from the group mean in the dimensions where the executed action lies more than one standard deviation from it, and toward it in the others. The two terms are the natural-gradient steps on the mean and the log standard deviation of a diagonal Gaussian in separable natural evolution strategies [91], with the advantage as the utility of the executed action and the mean and the spread estimated from the group.

Update. The directions are not rescaled per group. The regression is made at one step of the cached trajectory of each sample, drawn uniformly, toward $v_\mathrm{old} - \eta\,g_k - (1-t)\,J\,\eta\,g_k$, where J is the Jacobian of v_old in the sampler state at that step. The last term is the first-order response of the velocity when the whole trajectory is displaced so that it ends at $x_k + \eta\,g_k$ (Sec. 3.5), and it costs one Jacobian-vector product. The loss is the squared error divided, per action dimension, by the mean squared velocity change that is asked for, without the noise-level weight and the trust region. After each iteration, the samples at 2,048 held-out states are drawn again from the updated policy with the same noise, and η, which starts at 1, is multiplied by a factor between 0.8 and 1.25 so that their root-mean-square change approaches a target fraction of the spread of the samples. FlowRDM replaces the update of the released FPO++ algorithm with this update and keeps the networks, the 48 steps per iteration, the 16 epochs of 6 minibatches, and the perturbation of the released configuration. Its hyperparameters for this task are a coefficient of $1/16$ in place of $1/2$ in the second term of Eq. (10), the generalized advantage estimate with $\lambda = 0.9$, standardized over the iteration and clipped at ± 1, as the reward A, a target of 0.06 of the spread of the samples for the change at the held-out states, and a learning rate of 10^{-4}.

Baselines. All three methods use 4,096 environments and three seeds each. FPO++ is the released training script with its configuration for this task: 48 steps per iteration, an actor and a critic with three layers of 1,024, 512, and 256 units, a 64-step Euler sampler, 16 noise and time samples per action, 16 epochs of 6 minibatches, AdamW with an adaptive learning rate that starts at 10^{-4}, a clip parameter of 0.01, and a Gaussian perturbation of standard deviation 0.1 on the sampled actions. PPO is the Gaussian PPO [75] of RSL-RL [76] with the configuration that BeyondMimic releases for this task, unchanged: 24 steps per iteration, networks with three layers of 512, 256, and 128 units, an entropy coefficient of 0.005, a clip parameter of 0.2, 5 epochs of 4 minibatches, and a learning rate adapted to a KL divergence of 0.01. On Fall and get up, two of the three FPO++ seeds diverge, after 137M and 236M steps; Fig. 5 and Tab. 11 give the mean of the other two seeds from 137M steps on and the third seed alone from 236M steps on.

Evaluation. The training return is the mean return of the last 100 finished training episodes, with exploration noise and with start phases from the sampler of the task, and Fig. 5 averages it over 21 iterations (41 of PPO). The exploration noise of the Gaussian policy and of the flow policies differs, so we also compare the methods in an evaluation. It runs one episode in each of 4,096 environments, from start phases drawn uniformly over the motion with a fixed seed, the same for every checkpoint and method, with the sampler started from zero noise or with the mean action of PPO, and with the observation noise and the pushes of the task.

Motion	PPO	FPO++	FlowRDM (ours)	
Environment steps	295M	295M	295M	197M
Walk	41.8	42.7	43.8	43.2
Run	35.9	37.2	38.1	37.4
Dance 1	26.7	26.3	30.1	28.3
Dance 2	32.0	33.3	36.9	35.5
Fight	28.3	27.6	33.0	31.0
Jumps	35.0	36.0	38.2	37.3
Mean of the six above	33.3	33.9	36.7	35.5
Fall and get up	19.1	10.6[†]	30.2	26.5
Mean of all seven	31.3	–	35.7	34.2

Table 11. **Humanoid motion tracking from scratch** on the Unitree G1: return of the evaluation with zero-noise sampling on 4,096 start phases, mean of three seeds, after 295M environment steps and, for FlowRDM, also after 197M. [†]One seed; the other two diverge. PPO and FPO++ are our runs of the released configurations.

Results. Table 11 gives this evaluation after 295M steps, 1,500 iterations of FlowRDM and FPO++ and 3,000 of PPO. FlowRDM has the highest return on all seven motions, and its lowest seed is above every seed of both baselines on each motion. On the six motions of the FPO++ paper, 94% of its episodes reach the time limit, against 89% for PPO and 86% for FPO++, and its mean training return reaches the value that FPO++ has after 295M steps within 140M to 180M steps, depending on the motion, so it needs 1.6 to 2.1 times fewer steps (Fig. 5).

F.6. Details of Figure 1

Text-to-image. The prompt is one of Fig. 7. The curves show total reward on SD3.5-Medium over the 128-GPU-hour training budget of FlowRDM; longer runs are truncated by linear interpolation, and the reward axis shows the post-training range from 2.9 to 3.3. Labels beside the curves give peak gains over the base model's total reward of 2.0136: 53.9% for DiffusionART (ART, peak 3.0996), 55.0% for DiffusionNFT (NFT, 3.1211), and 59.7% for FlowRDM (3.2149). The corresponding gain on FLUX.1-dev is 14.1%, from 2.8031 to 3.1983 (Sec. 5.1).

World model. The item is a ball that rolls across a table, shown at 0, 0.25, 0.5, and 0.75 s of the continuation, from Cosmos3-Super and its post-trained model. The chart compares post-trained Cosmos3-Super with the four highest-scoring public entries in the video-to-video part of the Physics-IQ Verified leaderboard [63] (snapshot of October 20, 2026), sorted by score. Cosmos3 and Odyssey use best-of-eight sampling; FLUX 3 [large] reports best-of-N selection.

Closed-loop driving. The frame is from the front camera in the rollout of the base policy, 7.8 s after the start of a held-out scenario that the base policy fails and FlowRDM passes in the runs of Tab. 4, with the ground tracks of the two closed-loop rollouts over the following 4.5 s projected into it. The ring marks where the base policy leaves the road, 1.5 s after the frame; the post-trained policy drives on without an at-fault event. The bars, which start at zero, show kilometres per at-fault event from Tab. 4, including public Alpamayo 1.5, in descending order; GRPO denotes Flow-GRPO. The dashed line marks the base policy's 1.889 kilometres per at-fault event; FlowRDM's 2.749 is the highlighted gain of 45.5%.

Robot control. The stills show the Franka Panda arm on Square and the Unitree G1 on the motion Dance 2, each under a policy trained with FlowRDM. The manipulation curves show success rate on Square: FlowRDM is the mean of three seeds, and FPO++, FPO, and both DPPO implementations are the curves reported by the FPO++ authors [99]. The tracking curves show training return on Dance 2, averaged over three seeds, with the temporal averaging of Fig. 5. The brackets compare the environment steps needed to reach FPO++'s final displayed performance, with crossings linearly interpolated on the mean curves: FlowRDM reaches the success rate of 0.582 after 3.2M steps, against 8.0M for FPO++ (2.5×), and the training return of 24.8 after 155M steps, against 295M (1.9×).