

Finite Difference Flow Optimization for RL Post-Training of Text-to-Image Models

David McAllister*
UC Berkeley

Miika Aittala
NVIDIA

Tero Karras
NVIDIA

Janne Hellsten
NVIDIA

Angjoo Kanazawa
UC Berkeley

Timo Aila
NVIDIA

Samuli Laine
NVIDIA

Abstract

Reinforcement learning (RL) has become a standard technique for post-training diffusion-based image synthesis models, as it enables learning from reward signals to explicitly improve desirable aspects such as image quality and prompt alignment. In this paper, we propose an online RL variant that reduces the variance in the model updates by sampling paired trajectories and pulling the flow velocity in the direction of the more favorable image. Unlike existing methods that treat each sampling step as a separate policy action, we consider the entire sampling process as a single action. We experiment with both high-quality vision language models and off-the-shelf quality metrics for rewards, and evaluate the outputs using a broad set of metrics. Our method converges faster and yields higher output quality and prompt alignment than previous approaches.

1. Introduction

Generation of synthetic images based on text prompts has become a ubiquitous task for deep learning models. The dominant paradigm today is using diffusion models [21, 52, 53, 55, 56] that generate images by reversing a stochastic corruption process. These models learn a prompt-conditional probability flow that maps a random sample of noise into an image on the data manifold. Flow matching [2, 11, 19, 32, 35] is a popular parameterization for this flow [1].

As with large language models, the training of image synthesis models is commonly divided into pre-training and post-training stages that have different goals [43]. Pre-training is performed on large-scale, weakly curated data. The goal is to inject as much real-world knowledge into the model as possible with relatively little concern for image or prompt quality. In the typically much shorter post-training phase, the model’s sample distribution is then concentrated on regions

representing desirable outputs based on an external reward metric or a carefully curated fine-tuning set.

The post-training of image generators typically has a wide variety of vague goals, such as making the images more beautiful, have nicer composition, more expressive lighting, etc. This “true reward” cannot be expressed as a loss function, and thus cannot be directly optimized. What can be done instead is to define a set of quantifiable *proxy rewards* that hopefully address many of the desirable aspects. Reinforcement learning (RL) is a general method that can, in principle, optimize any such proxy reward. Common proxy rewards include learned models trained to mimic human preferences [27, 61] along with more targeted goals, such as the clarity of text rendering.

The pre-training objective of diffusion models is minimized by a unique flow field determined by the training dataset. As such, the training constantly supervises that the generated distribution stays in alignment. Unfortunately, this is no longer true during post-training, and aspects that are irrelevant to the specified reward are left to drift freely; for example, a reward targeting accurate text rendering can disregard or even harm the overall image quality. While this “reward hacking” [3] is in the nature of RL — the policy is only designed to optimize the provided reward — the RL algorithm should not exacerbate the problem by, e.g., actively drifting in random directions along the underspecified dimensions. Such bad behavior can be partially mitigated by, e.g., specifying multiple rewards, by limiting the amount of allowed change via KL regularization, or by mixing in some pre-training objective.

Existing RL post-training methods for diffusion models typically recast stochastic sampling as a Markov decision process (MDP), i.e., a sequence of policy-guided actions with Gaussian transitions [6, 16]. DDPO [6] adopts the PPO [51] algorithm for reward optimization under the MDP formulation. Flow-GRPO [34] and DanceGRPO [66] apply the same idea in the context of flow matching. In this approach, the proposed updates to the flow are random perturbations

*Work done during an internship at NVIDIA.

that are independent between the sampling steps. These perturbations are reinforced if their overall effect on the sampling trajectory lands on a higher-than-average reward.

The noise in these updates is a serious weakness from the viewpoint just laid out. While the aggregate update will improve the reward, only a small fraction of its magnitude contributes to this, while the rest is reward-neutral noise that pushes the flow around in random directions. This poses several problems. The speed of progress per update is significantly limited, as a large part of any individual update does not contribute to the goal. The noise also causes the unrelated dimensions to drift freely, e.g., cycling through random image styles if they are not constrained by any reward. Finally, the drift also slowly accumulates into detrimental side effects, essentially setting a cap on how much fine-tuning can be done in total. As an example of this, we show that Flow-GRPO starts to introduce artifacts into the generated images upon extended training.

The goal of our approach is to improve the signal-to-noise ratio of the flow updates. Our method uses similar random perturbations to discover high-reward images, but decouples the resulting update from this random walk. Specifically, our method generates two nearby images and uses their difference as an approximate gradient. This image difference is weighted by the respective reward difference, so that it is guaranteed to point from the lower-reward image to the higher-reward one. We then *uniformly* update the flow directions along the generating trajectory towards this direction, relying on the de-facto “non-rotational” behavior specific to diffusion flows [25]. Each sampling step thus receives an update that directly benefits the reward. This is in contrast to the MDP formulation where, roughly speaking, close to half of the individual flow updates may be detrimental to the reward.

Compared to Flow-GRPO, our method converges significantly faster, to higher rewards, and with fewer reward hacking artifacts. It can be used as a drop-in replacement for SOTA RL algorithms in post-training of diffusion models. These results position our formulation as a promising alternative for RL post-training.

Our implementation and trained models are available at <https://github.com/NVlabs/finite-difference-flow-optimization>

2. Previous Work

Many methods have been proposed for post-training neural networks, and we will briefly review the most closely related ones here. For a broader context, Liu et al. [33] present a survey on image generator post-training, constructing a taxonomy for a wide range of literature on related algorithms and evaluation. Also, Uehara et al. [57] present a more RL-focused survey that examines the connections between different post-training methods in detail.

Supervised fine-tuning A common approach is to simply continue training using a smaller fine-tuning dataset. Direct preference optimization (DPO) [48], on the other hand, relies on annotated human preference pairs. It is simple to implement, building on supervised training without needing an online reward model. However, it is limited to cases where annotated preference data is available. Diffusion DPO [58] applies the method to image generators.

Differentiable rewards If the reward function is differentiable with respect to the generated image, it is possible to backpropagate the reward gradients to the generator. Commonly used approaches in this category include ReFL [64], DRaFT [10], and Deep Reward Supervision (DRS) [62]. These methods differ mainly in where exactly the supervision happens inside the generator — ReFL supervises on a random step, DRaFT on a sequence of last steps, and DRS on a sparse subset of steps. Dual-process image generation [38] propagates gradients through a VLM reward to adjust image generator weights at inference time.

Domingo-Enrich et al. [12] cast reward fine-tuning as a stochastic optimal control (SOC) problem, which is strongly connected to RL [30]. They formulate a noise schedule that is free of value function bias that typical post-training objectives suffer from. They also present adjoint matching, a general method for solving SOC problems that removes high-variance importance weights from the regression objective.

Non-differentiable rewards In many practically relevant scenarios, the reward function is not differentiable, most notably when humans provide direct reward signals. Although this is an obvious application of RL, a number of RL-adjacent supervised learning approaches have been proposed as well. Off-policy sample-evaluate-update loops are used in language modeling [42, 67] and control settings [45, 49]. In image generation, several authors apply these techniques through reward-weighted regression [6, 15, 29] as well as rejection sampling [13].

True RL methods, applicable to any reward, include the aforementioned MDP-based methods DDPO [6], Flow-GRPO [34] and DanceGRPO [66], with the last two demonstrating current state-of-the-art results. Other works perform on-policy RL with the standard denoising objective [39, 65], though they rely on strong regularization for stability in image generation settings. Some works also apply variants of value-based RL [60] to diffusion models [23, 44, 46].

GFlowNets [5, 36, 68] also rely on the MDP formulation but propose a novel training objective that is closely related to path consistency learning [41, 57]. The aim is to tilt the distribution proportionally to the reward, the same goal as KL-regularized RL [30]. However, performance vs. mainstream SOTA (Flow-GRPO, DanceGRPO) is unknown.

Inference-time optimization Orthogonal to model post-training, inference-time tweaks to the sampling process can




Figure 1. Illustration of differences between denoising MDPs, including Flow-GRPO [34], and our method. **(a)** For each prompt, Flow-GRPO samples a group of trajectories using stochastic Euler–Maruyama sampling. The relative advantage A of each sample is the reward normalized with the group statistics. For each time step, the flow velocity is optimized towards the stochastic perturbation the trajectory took, or away from it if the sample’s advantage is negative. **(b)** In our method, we sample a pair of trajectories and compute the difference $\Delta \mathbf{x}$ between the output images. The flow velocity is optimized towards normalized $\Delta \mathbf{x}$ scaled by the reward delta ΔR . Note that both Flow-GRPO and our method apply updates to all sampled trajectories, whereas only one is highlighted here for clarity.

provide improvements to output quality at the cost of reducing diversity. These techniques include, e.g., classifier-free guidance (CFG) [20], gradient-guided diffusion [18], and loss-guided diffusion [54].

3. Background

Flow matching [2, 11, 19, 32, 35] is a parameterization of a probability flow diffusion process [21, 52, 53, 55, 56]. It is popular for its intuitive formulation and ease of implementation, as well as good practical performance stemming from favorable built-in design choices.

Flow matching works by drawing an initial random noise image $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_N)$ of N pixels, and numerically solving an ordinary differential equation (ODE)

$$d\mathbf{x}(t) = v_\theta(\mathbf{x}; t, \mathbf{c})dt \quad (1)$$

backward in time from $t = 1$ to $t = 0$ with initial condition $\mathbf{x}(1) = \epsilon$. Here, $v_\theta(\cdot)$ is the *velocity function*, implemented as a neural network with weights θ , that points towards the average of noise-free images consistent with the noisy image $\mathbf{x}(t)$ and time parameter t , conditioned on prompt embedding $\mathbf{c}$ for conditional generation. Following the ODE pushes the image towards the evolving denoising estimate under the training image dataset, gradually reducing the noise level and revealing a clean image $\mathbf{x}(0)$ from this distribution.

In practice, the ODE is solved by evaluating $\mathbf{x}(t)$ at discrete time steps $t_0 = 1, t_1, \dots, t_{T-1}, t_T = 0$, where each step yields an intermediate image $\mathbf{x}_i := \mathbf{x}(t_i)$ that is slightly less noisy than the previous one. Using the simple Euler solver scheme, the step from time t_i to t_{i+1} is given by $\mathbf{x}_{i+1} = \mathbf{x}_i + (t_{i+1} - t_i)v_\theta(\mathbf{x}_i; t_i, \mathbf{c})$, so that after T steps the initial noise image at t_0 is transformed into the generated im-

age at t_T . We refer to a sequence of images $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_T$ as a sampling trajectory.

3.1. Reward Maximization

Given a reward function $R(\mathbf{x})$ that maps images to scalar reward values, our goal is to fine-tune the pre-trained velocity function v_θ to maximize the expected reward over draws from the deterministically sampled generative model:

$$\arg \max_\theta \mathbb{E}_{\mathbf{c} \sim C, \mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} R(f_\theta(\mathbf{x}_0; \mathbf{c})), \quad (2)$$

where C is the distribution of fine-tuning prompt embeddings, and $f_\theta(\cdot)$ represents the process of drawing a sample starting from an initial noise $\mathbf{x}_0$ using velocity function v_θ . Geometrically, the goal is to redirect the flow velocities such that a larger mass of noises map to high-reward regions of the image space. Generally, the reward may be non-differentiable or highly discontinuous.

3.2. Markov Decision Process Formulation

As an alternative to ODE-based deterministic sampling, diffusion models also allow for SDE-based *stochastic* sampling under the same flow. Stochastic sampling injects fresh Gaussian noise at each time step to randomize the trajectory.

MDP based approaches [6, 16, 34, 66] take advantage of stochasticity and cast each stochastic step as a Gaussian-distributed action distribution whose mean is defined by the flow velocity. Beneficial actions are reinforced by pulling the velocity toward the random steps of high-reward trajectories.

Fig. 1a illustrates the group-relative approach used by Flow-GRPO [34] and DanceGRPO [66] algorithms. They adopt the denoising MDP for flow schedules and estimate the advantage of each trajectory by normalizing its corresponding reward within a group. The advantages determine the

weight of the corresponding updates to the flow. While these updates point towards more favorable images on average, a significant proportion of them point opposite of reward increase, contributing to undesirable variance.

4. Our Method

We adopt an approximate on-policy approach of alternating between (1) generating a set of trajectory rollouts from the current model, and (2) training the network velocity predictions along these trajectories, so as to redirect them towards collecting higher reward at their endpoints. To obtain a training signal without direct access to gradient of the reward, we propose to roll out *pairs* of images with variations in detail and to reinforce the probability of the more favorable image among the two.

Specifically, we generate a pair of rollouts $\mathbf{x}_i$ and $\hat{\mathbf{x}}_i$ from a shared initial noise $\mathbf{x}_0 = \hat{\mathbf{x}}_0 = \epsilon$, but apply a modest amount of stochasticity that perturbs them along the sampling trajectory. This induces random differences in generated image details, whereby one of the output images $\{\mathbf{x}_T, \hat{\mathbf{x}}_T\}$ usually collects a higher reward than the other. Then, the difference vector $\Delta\mathbf{x} = \hat{\mathbf{x}}_T - \mathbf{x}_T$ weighted by the reward difference $\Delta R = R(\hat{\mathbf{x}}_T) - R(\mathbf{x}_T)$ will point towards the higher-reward image. We train the velocities at all time steps along both trajectories to bend towards $\Delta R \Delta\mathbf{x}$. The approach is illustrated in Fig. 1b.

4.1. Analysis

The weighted difference $\Delta R \Delta\mathbf{x}$ points towards desirable changes in the generated image at $t = 0$, whereas our updates divert the flow toward this direction at intermediate steps. Our method thus relies on the expectation that this induces a similar change in the image.

To justify this informally, we note that the denoising process essentially fades out noise to reveal a hidden signal in a coarse-to-fine fashion. Then, adding “signal”, i.e., the reinforced image difference, at an intermediate noisy image will approximately pass it to the generated image. This is a core underlying assumption in a broad range of methods (e.g., [8, 9, 26, 40]) that perform coarse edits on a partially noised image (say, adding a blob of green pixels), and rely on the remainder of the flow to flesh out the detail while maintaining the broad direction of the edit (say, into a tree).

More formally, our update aims to satisfy $\nabla R(\mathbf{x}_T)^\top \mathbf{J}_i(\mathbf{x}_i) [\Delta R \Delta\mathbf{x}] \geq 0$ for each step i , where $\mathbf{J}_i(\mathbf{x}_i)$ is the Jacobian matrix of the mapping induced by the flow from step i onward. While it is possible to construct counter-examples that violate this condition, there are significant arguments for it holding at least to the extent required in practical image generation. For example, it can be shown that under certain assumptions the condition holds on expectation if $\mathbf{J}_i$ is positive semi-definite (see Appendix C), which would be true for an optimal transport

Algorithm 1 Stochastic flow matching sampler

```

procedure STOCHASTICFLOWAMPLER( $\theta, \epsilon, \mathbf{c}, \gamma_{0:T}$ )
   $\mathbf{x}_0 \leftarrow \epsilon$                                 ▷ Start from given noise
  for each  $i \in \{0, \dots, T-1\}$  do             ▷ Take  $T$  sampling steps
     $\tilde{t}_{i+1} \leftarrow t_{i+1} / (1 - \gamma_i t_{i+1} + \gamma_i)$   ▷ Overshoot
     $\mathbf{v}_i \leftarrow v_\theta(\mathbf{x}_i; t_i, \mathbf{c})$            ▷ Evaluate velocity
     $\tilde{\mathbf{x}}_{i+1} \leftarrow \mathbf{x}_i + (\tilde{t}_{i+1} - t_i) \mathbf{v}_i$     ▷ Step from  $t_i$  to  $\tilde{t}_{i+1}$ 
    sample  $\epsilon_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$                 ▷ Draw fresh noise
     $\tilde{\mathbf{x}}_{i+1} \leftarrow \frac{\tilde{\mathbf{x}}_{i+1} + \tilde{t}_{i+1} \sqrt{\gamma_i^2 + 2\gamma_i} \cdot \epsilon_i}{\gamma_i \tilde{t}_{i+1} + 1}$   ▷ Add noise to land
     $\mathbf{x}_{i+1} \leftarrow$                                at  $t_{i+1}$ , correct scale
  return  $\mathbf{x}_{0:T}, \mathbf{v}_{0:T-1}$ 

```

mapping [50]. The apparent similarity of diffusion flows and optimal transport mappings, while not theoretically exact [28], has inspired studies where a strong numerical similarity is nevertheless found [25].

4.2. Stochastic Sampling

Like Flow-GRPO, our approach requires a means to generate nearby random image variants. As discussed by Karras et al. [24], an Euler–Maruyama sampler (used in, e.g., Flow-GRPO) suffers from two related numerical problems when applied to flows. The velocity network is called with a slightly inconsistent time conditioning and noise amount at each step, as the impact of the noise injection sub-step is not accounted for. This is exacerbated by oversized noise injections that occur at some steps, because the amount of noise added is not proportioned to the existing noise level in the sample.

In Algorithm 1 we adapt the key idea from the EDM stochastic sampler [24] to introduce stochasticity into the flow-matching solution trajectories. When stepping from t_i to t_{i+1} , we take the ODE direction but overshoot the target time to a lower noise level, and then compensate by adding fresh random noise to land at t_{i+1} (with appropriate scale corrections, see Appendix B.1 for details).

The parameters γ_i specify the strength of stochastic randomization at each solver step. Roughly, γ_i expresses the fraction of noise in $\mathbf{x}_i$ that is re-randomized on that step, with value $\gamma_i = 0$ corresponding to deterministic flow matching. Intuitively, re-randomizing at an intermediate noise level replaces the to-be-generated finer image detail that is still encoded by the noise with a different random realization, while keeping the already resolved coarser detail intact.

4.3. Implementation Details

We present the full method pseudocode in Appendix B.

Stochasticity strength Where not specified otherwise, we use a uniform stochasticity schedule of $\gamma_i = 0.0025$ for all timesteps. This induces variations in semantic detail of an image pair (e.g., moving or swapping out parts), but mostly retains their overall layout and content. We experimented

with more complex schedules but found no consistent benefit to them. See Appendix B.2 for details.

Normalization Neural network training benefits from approximately uniform gradient magnitude among training samples. The random stochastic perturbations result in different magnitudes for $\Delta \mathbf{x}$ between rollouts, and we can also expect ΔR to be approximately (to the first order) proportional to $\Delta \mathbf{x}$. The raw training signal $\Delta R \Delta \mathbf{x}$ thus counts the norm $\|\Delta \mathbf{x}\|_{\text{RMS}}$ twice into its magnitude. We cancel this effect by using normalized $\Delta \mathbf{x} = \Delta \mathbf{x} / (\|\Delta \mathbf{x}\|_{\text{RMS}}^2 + 10^{-6})$ as the actual training signal.

Batching We train the model for up to 1000 epochs. For each epoch, we generate 432 pairs of rollouts with randomly drawn prompts and store all $T = 40$ steps of their trajectories, resulting in the same memory consumption as Flow-GRPO. The resulting $432 \times 2 \times 40$ samples are then randomly split into 4 training batches, totaling 8640 samples/batch. The weighted image differences are backpropagated through the velocity function and accumulated, until they are updated into model weights θ after each batch using AdamW [37].

On-policy optimization Making several training steps on a frozen set of rollouts throughout the epoch is sample-efficient, as computing fresh rollouts frequently would be expensive. However, it runs the risk of training on stale data, as the parameters θ may have changed significantly since the set of rollouts was refreshed. Following the usual RL policy optimization recipe, we apply clipping to downweight updates to velocities that have moved too far from their original position since the last refresh. We use Simple Policy Optimization (SPO) [63] clipping that is similar to the widely used Proximal Policy Optimization (PPO-Clip) [51].

5. Results

We will now evaluate our proposed RL algorithm experimentally to compare its performance against Flow-GRPO [34], a current state-of-the-art method. We start by considering the effectiveness of our method in increasing the proxy reward as quickly as possible, using rewards of different complexity to highlight differences between the algorithms. In these tests, we are mainly interested in convergence speed, highest obtainable reward, and potential algorithm-dependent artifacts.

We then employ external control metrics to analyze how the rewards, RL algorithms, and CFG strength affect image quality, prompt alignment, and diversity. Finally, we ablate the major design choices of our algorithm as well as the key differences to Flow-GRPO. Additional results, metrics, and details for all experiments are provided in Appendix A.

5.1. Experiment Setup

We use the official implementation of Flow-GRPO [34]. We explored its hyperparameters to a considerable extent and

concluded that the official defaults are close to optimal in all cases. We thus use them with the following exceptions. First, we disable exponential moving averaging of model weights as it was strictly harmful. Second, unless stated otherwise, we disable KL regularization [34] and CFG [20] because they can mask the differences between RL algorithms and ideally would not be needed. Third, we use 40 sampling steps during both training and inference. The role of CFG is explored in Section 5.3, and reducing the number of sampling steps during training is evaluated in Section 5.5.

We perform all experiments using Stable Diffusion 3.5 Medium [14], extended with low-rank adaptation layers (LoRA) [22] and update only these layers during post-training, in line with Flow-GRPO. We generate all images at 512×512 resolution. The number of model and reward evaluations per epoch is the same for Flow-GRPO and our method.

We use two primary reward functions, PickScore[47] and a novel VLM reward, that focus on predicting human preferences and prompt alignment, respectively. Both are evaluated using the Pick-a-Pic [27] training prompts. Our use of a VLM reward is motivated by the observation made by Lin et al. [31] that VLMs are surprisingly effective in estimating things like prompt alignment compared to explicit human preference models, and their performance can be expected to further improve as new models are released. To this end, we adopt a scheme similar to VQAScore [31] using Qwen2.5-VL-7B-Instruct [59]. We feed in the generated image along with a text query “Does this image match the caption “...”? Answer Yes or No.”, where ... corresponds to the prompt that was used to generate the image. We then run the VLM to predict the logits for the next token and calculate the reward as $100 \cdot \text{sigmoid}(\text{logits}[\text{Yes}] - \text{logits}[\text{No}])$. We also consider the weighted combination of these two rewards, defined as $\text{PickScore} + \text{VLM alignment} / 10$.

For evaluation, we use the reward itself as well as a set of external control metrics, including OneIG-Bench [7] and HPSv2 [61]. We calculate these after-the-fact based on checkpoints exported during training, so that the evaluation protocol remains independent of the post-training algorithm.

5.2. Reward Optimization

Fig. 2 compares our method with Flow-GRPO when optimizing three different rewards: (a) PickScore, (b) VLM-based prompt alignment, and (c) their weighted combination. In this test, an epoch is equally expensive for both methods. PickScore is the easiest to optimize and both methods work well, although our method reaches a slightly higher reward much faster. The result achieved by Flow-GRPO matches the original paper closely.

The VLM-based alignment reward is more difficult to optimize; Flow-GRPO starts to struggle and the gap to our method widens significantly. The combined reward paints




Figure 2. Comparison between our method and Flow-GRPO using different reward functions. Each plot shows the reward as a function of RL epoch. KL regularization and CFG are disabled. Example images for the highlighted checkpoints of combined reward are shown visually in Fig. 3. **(a)** PickScore reward. Flow-GRPO reaches 23.43 after 1000 epochs (2000 training steps). **(b)** Our VLM alignment reward: *Does this image match the caption "..."? Answer Yes or No.* **(c)** Combined reward: PickScore + VLM alignment / 10.



Figure 3. Visual results for the combined reward using the prompt “A girl with pigtails is holding a giant sunflower.” The epoch numbers (bottom-right corners) correspond to the dots in Fig. 2c, approximately matching the reward between the methods. The last column: Flow-GRPO is unable to increase the reward enough to match our method.



Figure 4. When trained long enough, Flow-GRPO starts exhibiting grid-like artifacts that come and go, but are very disturbing in some epochs (here 830). Our results do not show these at any stage; here epoch 70 serves as an iso-reward comparison. Left: “A green tractor plowing a field at sunrise.” Right: “A quaint cottage nestled in a vibrant flower-filled meadow.” Combined reward, no CFG or KL regularization.




Figure 5. Control metrics for our three rewards (combined, VLM, PickScore) as a function of RL epoch. **(a)** Independent assessment of prompt alignment. **(b)** Independent assessment of human preferences. **(c)** Independent assessment of result diversity.




Figure 6. Control metrics without CFG ($w = 1$), and with medium ($w = 2$) and high ($w = 4.5$) guidance strength as a function of RL epoch using the combined reward. CFG drastically changes the results before post-training (epoch 0) but the gap narrows over time. Results for the highlighted checkpoints are shown visually in Figure 7.




Figure 7. Visual results for prompt alignment vs. diversity: “A cat wearing ski goggles is exploring in the snow.” **(a)** Without CFG, the original model has high diversity but poor quality and alignment. **(b)** Enabling CFG flips all these axes, leading to a distinctive low-detail look. **(c)** Our RL post-training achieves similar diversity and alignment with more detail. **(d)** Enabling CFG simplifies images and reduces diversity slightly. The used snapshots align with the dots in Fig. 6.




Figure 8. OneIG-Bench alignment score for ablations of our method using the combined reward. The plots correspond to the top and bottom section of the table, respectively. We report the alignment score at 200 epochs as well as the highest achieved score, annotated in the plots.

a similar picture with our method converging significantly faster to a higher reward. More complicated rewards consisting of multiple goals can be beneficial as they can reduce the aspects the RL optimization is blind to.

Fig. 3 shows representative image progressions for the combined reward. The images are drawn from training snapshots corresponding to the dots annotated in Fig. 2 and demonstrate that higher rewards correspond to visually better results, as expected. While subjective comparison is inherently difficult, we feel that our method stacks well against Flow-GRPO in reward-matched pairwise comparison. Furthermore, there is a clear style drift in the Flow-GRPO images — this was a consistent behavior with all the prompts we experimented with. We attribute this style drift to Flow-GRPO’s more noisy flow updates that lead to a greater amount of random mutation of the flow until a given reward is reached.

We also note that Flow-GRPO eventually starts introducing grid-like reward hacking artifacts that fade in and out during training. Fig. 4 shows handpicked examples where this effect is severe. The fact that these artifacts again disappear over time implies that the combined reward likely does penalize them, but not very strongly. We have never seen these artifacts in our method, even after equally long training.

5.3. Control Metrics

Let us now further analyze the post-trained models to see how the choice of reward affects alignment with prompts, alignment with human preferences, and diversity. To estimate these, we use OneIG-Bench [7] (prompt alignment, diversity) and HPSv2 [61] (human preference) as independent control metrics.

Fig. 5a shows that the prompt alignment control metric is poorly optimized by the PickScore reward alone, as could be expected. Our VLM reward, which specifically targets this aspect, does a much better job, but the combined reward works even more reliably. Same conclusions apply to both our algorithm and Flow-GRPO, and the benefits of our

method that were observed with the raw reward (Fig. 2b) also hold for the control metric. Fig. 5b shows that human preferences are well optimized by PickScore that focuses on this aspect. VLM reward does not work particularly well here, but again the combined reward comes out on top using our algorithm. Fig. 5c shows that all rewards reduce diversity, with PickScore causing the fastest decline and VLM reward the slowest. Considering that our algorithm optimizes the rewards much faster than Flow-GRPO, diversity loss is approximately equally fast for the two algorithms on an iso-reward basis. These observations position the combined reward as a clear choice over the more commonly used PickScore.

Fig. 6 uses the same control metrics to assess the effects of CFG. In these tests, we enable CFG for the base model and redo the entire RL post-training and evaluation using the CFG-adjusted velocity predictions in place of the original ones. We can see that increasing the guidance strength improves alignment with prompts and human preferences, at the cost of further diversity loss. This is not surprising, and the observations apply equally to our method and Flow-GRPO.

Fig. 7 provides visual examples of this tradeoff. Before post-training, the image quality of the base model is poor without CFG (Fig. 7a). Enabling CFG in that setting (Fig. 7b) improves prompt alignment and image quality dramatically while reducing diversity (see Fig. 6abc, epoch 0). With post-training, it is debatable whether CFG is useful, especially given its high inference cost. Subjective inspection of the result images suggests that post-training without CFG improves image fidelity and detail significantly (Fig. 7c), and the use of CFG in this setting mainly compromises diversity and introduces the characteristic, high-contrast look (Fig. 7d).

5.4. Ablations

Figure 8 presents a number of ablation experiments, starting with the Flow-GRPO baseline (row A). The amount of stochasticity is tuned separately for each ablation vari-





Figure 9. Wall-clock time comparison (NVIDIA H200 hours). (a) Actual elapsed time. (b) Assuming ideal, zero-overhead implementation of both methods. Note the different horizontal scales.

ant. Switching to our stochastic sampler in Algorithm 1 (row B) improves convergence considerably. Replacing the Flow-GRPO policy gradients with our finite difference flow optimization (row C), while keeping the number of unique prompts per epoch unchanged, results in another major boost. We then start all trajectories in each group from the same initial noise (row D), which improves stability and improves the results further. Finally, we can increase the number of unique prompts considered in each epoch by virtue of needing only two rollouts per prompt, which leads to our full method (row E).

The bottom section of the table evaluates some of our specific design choices. Instead of using stochastic sampling for both trajectories (row E), we can equally well use deterministic sampling for one of them (row F), which may be desirable in order to reduce the discrepancy between rollout generation and inference-time sampling. Using PPO [51] (row G) in place of SPO [63] (row E) yields similar results, although only when using the interval stochasticity schedule (Appendix B.2)—this was the only experiment where a constant γ_i schedule was clearly inferior to a more complex schedule. It is beneficial to normalize Δx as described in Section 4.3; disabling this (row H) is highly detrimental to stability. Replacing $\Delta R \Delta x$ with the actual reward gradient obtained by backpropagating through the reward model (row I) does not improve the results, implying that our method is a good fit regardless of whether the reward is differentiable or not. Further backpropagating this reward gradient through the sampling trajectory to determine the update for each sampling step (row J) is significantly worse than our approach of bending all sampling steps towards $\Delta R \Delta x$.

5.5. Wall-Clock Performance

Fig. 9a plots convergence as a function of GPU hours, comparing our method and Flow-GRPO in baseline (40 sampling steps) and fast (10 steps) configurations. As can be seen, our method converges significantly faster, crossing the high-

lighted combined reward level $19\times$ faster in the baseline configuration and $5\times$ faster in the fast configuration. Because the official Flow-GRPO implementation suffers from significant implementation overheads, we also estimated zero-overhead convergence curves (Fig. 9b) to confirm that our wall-clock advantage is not merely a result of a more careful implementation.

6. Discussion and Future Work

Our method (FDFO) presents a step forward in RL post-training of diffusion models. It is a direct replacement to Flow-GRPO with faster convergence and higher-quality results. By leaving the multi-step MDP formulation, we shorten the reward attribution horizon and expand the algorithmic design space.

Our evaluation is focused on comparing RL algorithms with as few confounding factors as possible, and thus we disabled KL regularization in our main experiments. However, it remains compatible with our algorithm, as demonstrated in further experiments in Appendix A. KL regularization is typically used to better retain variation in the results, but the best way to do this in general remains an open problem. Finding a way to formulate a VLM-based reward component that targets diversity directly could be an interesting avenue for future work.

Acknowledgments We extend our thanks to Tero Kuosmanen and Samuel Klenberg for maintaining our compute infrastructure.

References

- [1] Michael Albergo, Nicholas Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *CoRR*, abs/2303.08797, 2023. 1
- [2] Michael S. Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. In *Proc. ICLR*, 2023. 1, 3

- [3] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in AI safety. *CoRR*, abs/1606.06565, 2016. 1
- [4] Uri M. Ascher and Linda R. Petzold. *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations*. Society for Industrial and Applied Mathematics, 1998. 21
- [5] Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. In *Proc. NeurIPS*, 2021. 2
- [6] Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion models with reinforcement learning. In *Proc. ICLR*, 2024. 1, 2, 3
- [7] Jingjing Chang, Yixiao Fang, Peng Xing, Shuhan Wu, Wei Cheng, Rui Wang, Xianfang Zeng, Gang Yu, and Hai-Bao Chen. OneIG-Bench: Omni-dimensional nuanced evaluation for image generation. In *Proc. NeurIPS*, 2025. 5, 8, 13
- [8] Sherry X. Chen, Yaron Vaxman, Elad Ben Baruch, David Asulin, Aviad Moreshet, Kuo-Chin Lien, Misha Sra, and Pradeep Sen. TiNO-Edit: Timestep and noise optimization for robust diffusion-based image editing. In *Proc. CVPR*, 2024. 4
- [9] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. ILVR: Conditioning method for denoising diffusion probabilistic models. In *Proc. ICCV*, 2021. 4
- [10] Kevin Clark, Paul Vicol, Kevin Swersky, and David J. Fleet. Directly fine-tuning diffusion models on differentiable rewards (DRaFT). In *Proc. ICLR*, 2024. 2
- [11] Mauricio Delbracio and Peyman Milanfar. Inversion by direct iteration: An alternative to denoising diffusion for image restoration. *TMLR*, 2023. 1, 3
- [12] Carles Domingo-Enrich, Michal Drozdal, Brian Karrer, and Ricky T. Q. Chen. Adjoint matching: Fine-tuning flow and diffusion generative models with memoryless stochastic optimal control. In *Proc. ICLR*, 2025. 2
- [13] 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. *TMLR*, 2023. 2
- [14] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis. In *Proc. ICML*, 2024. <https://huggingface.co/stabilityai/stable-diffusion-3.5-medium>. 5, 23
- [15] 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 *Proc. ICLR*, 2025. 2, 13
- [16] Ying Fan, Olivia Watkins, Yuqing Du, Hao Liu, Moonkyung Ryu, Craig Boutilier, Pieter Abbeel, Mohammad Ghavamzadeh, Kangwook Lee, and Kimin Lee. DPOR: Reinforcement learning for fine-tuning text-to-image diffusion models. In *Proc. NeurIPS*, 2023. 1, 3
- [17] Stephanie Fu, Netanel Tamir, Shobhita Sundaram, Lucy Chai, Richard Zhang, Tali Dekel, and Phillip Isola. DreamSim: Learning new dimensions of human visual similarity using synthetic data. In *Proc. NeurIPS*, 2023. 13
- [18] Yingqing Guo, Hui Yuan, Yukang Yang, Minshuo Chen, and Mengdi Wang. Gradient guidance for diffusion models: An optimization perspective. In *Proc. NeurIPS*, 2024. 3
- [19] Eric Heitz, Laurent Belcour, and Thomas Chambon. Iterative α -(de)blending: A minimalist deterministic diffusion model. In *Proc. SIGGRAPH*, 2023. 1, 3
- [20] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *Proc. NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. 3, 5
- [21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Proc. NeurIPS*, 2020. 1, 3, 21
- [22] 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 *Proc. ICLR*, 2022. 5
- [23] Zhiwei Jia, Yuesong Nan, Huixi Zhao, and Gengdai Liu. Reward fine-tuning two-step diffusion models via learning differentiable latent-space surrogate reward. In *Proc. CVPR*, 2025. 2
- [24] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Proc. NeurIPS*, 2022. 4, 20, 21
- [25] Valentin Khruikov, Gleb Ryzhakov, Andrei Chertkov, and Ivan Oseledets. Understanding DDPM latent codes through optimal transport. In *Proc. ICLR*, 2023. 2, 4
- [26] Gwanghyun Kim, Taesung Kwon, and Jong Chul Ye. DiffusionCLIP: Text-guided diffusion models for robust image manipulation. In *Proc. CVPR*, 2022. 4
- [27] 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 *Proc. NeurIPS*, 2023. https://huggingface.co/yuvalkirstain/PickScore_v1. 1, 5, 13, 23
- [28] Hugo Lavenant and Filippo Santambrogio. The flow map of the Fokker-Planck equation does not provide optimal transport. *Appl. Math. Lett.*, 133, 2022. 4
- [29] Kimin Lee, Hao Liu, Moonkyung Ryu, Olivia Watkins, Yuqing Du, Craig Boutilier, Pieter Abbeel, Mohammad Ghavamzadeh, and Shixiang Gu. Aligning text-to-image models using human feedback. *CoRR*, abs/2302.12192, 2023. 2
- [30] Sergey Levine. Reinforcement learning and control as probabilistic inference: Tutorial and review. *CoRR*, abs/1805.00909, 2018. 2
- [31] Zhiqiu Lin, Deepak Pathak, Baiqi Li, Jiayao Li, Xide Xia, Graham Neubig, Pengchuan Zhang, and Deva Ramanan. Evaluating text-to-visual generation with image-to-text generation. In *Proc. ECCV*, 2024. 5
- [32] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *Proc. ICLR*, 2023. 1, 3

- [33] Buhua Liu, Shitong Shao, Bao Li, Lichen Bai, Zhiqiang Xu, Haoyi Xiong, James T. Kwok, Sumi Helal, and Zeke Xie. Alignment of diffusion models: Fundamentals, challenges, and future. *ACM Comput. Surv.*, 2026. 2
- [34] Jie Liu, Gongye Liu, Jiajun Liang, Yangguang Li, Jiaheng Liu, Xintao Wang, Pengfei Wan, Di Zhang, and Wanli Ouyang. Flow-GRPO: Training flow matching models via online RL. In *Proc. NeurIPS*, 2025. https://github.com/yifan123/flow_grpo. 1, 2, 3, 5, 13, 21
- [35] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *Proc. ICLR*, 2023. 1, 3, 21
- [36] Zhen Liu, Tim Z. Xiao, Weiyang Liu, Yoshua Bengio, and Dinghui Zhang. Efficient diversity-preserving diffusion alignment via gradient-informed GFlowNets. In *Proc. ICLR*, 2025. 2
- [37] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Proc. ICLR*, 2019. 5
- [38] Grace Luo, Jonathan Granskog, Aleksander Holynski, and Trevor Darrell. Dual-process image generation. In *Proc. ICCV*, 2025. 2
- [39] David McAllister, Songwei Ge, Brent Yi, Chung Min Kim, Ethan Weber, Hongsuk Choi, Haiwen Feng, and Angjoo Kanazawa. Flow matching policy gradients. In *Proc. ICLR*, 2026. 2, 23
- [40] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. SDEdit: Guided image synthesis and editing with stochastic differential equations. In *Proc. ICLR*, 2022. 4
- [41] Ofir Nachum, Mohammad Norouzi, Kelvin Xu, and Dale Schuurmans. Bridging the gap between value and policy based reinforcement learning. In *Proc. NeurIPS*, 2017. 2
- [42] Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, Xu Jiang, Karl Cobbe, Tyna Eloundou, Gretchen Krueger, Kevin Button, Matthew Knight, Benjamin Chess, and John Schulman. WebGPT: Browser-assisted question-answering with human feedback. *CoRR*, abs/2112.09332, 2022. 2
- [43] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In *Proc. NeurIPS*, 2022. 1
- [44] Seohong Park, Qiyang Li, and Sergey Levine. Flow Q-learning. In *Proc. ICML*, 2025. 2
- [45] Jan Peters and Stefan Schaal. Reinforcement learning by reward-weighted regression for operational space control. In *Proc. ICML*, 2007. 2
- [46] Michael Psenka, Alejandro Escontrela, Pieter Abbeel, and Yi Ma. Learning a diffusion model policy from rewards via Q-score matching. In *Proc. ICML*, 2024. 2
- [47] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proc. ICML*, 2021. <https://huggingface.co/laion/CLIP-ViT-H-14-laion2B-s32B-b79K> and <https://huggingface.co/openai/clip-vit-large-patch14>. 5, 13
- [48] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *Proc. NeurIPS*, 2023. 2
- [49] Reuven Rubinstein. The cross-entropy method for combinatorial and continuous optimization. *Method. Comput. Appl. Prob.*, 1(2), 1999. 2
- [50] Filippo Santambrogio. *Optimal Transport for Applied Mathematicians: Calculus of Variations, PDEs, and Modeling*. Birkhäuser, 2015. 4
- [51] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347, 2017. 1, 5, 9
- [52] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proc. ICML*, 2015. 1, 3
- [53] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *Proc. ICLR*, 2021. 1, 3, 21
- [54] Jiaming Song, Qingsheng Zhang, Hongxu Yin, Morteza Mardani, Ming-Yu Liu, Jan Kautz, Yongxin Chen, and Arash Vahdat. Loss-guided diffusion models for plug-and-play controllable generation. In *Proc. ICML*, 2023. 3
- [55] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Proc. NeurIPS*, 2019. 1, 3
- [56] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *Proc. ICLR*, 2021. 1, 3, 20, 21
- [57] Masatoshi Uehara, Yulai Zhao, Tommaso Biancalani, and Sergey Levine. Understanding reinforcement learning-based fine-tuning of diffusion models: A tutorial and review. *CoRR*, abs/2407.13734, 2024. 2
- [58] 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 *Proc. CVPR*, 2024. 2
- [59] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. Qwen2-VL: Enhancing vision-language model’s perception of the world at any resolution. *CoRR*, abs/2409.12191, 2024. <https://huggingface.co/Qwen/Qwen2.5-VL-7B-Instruct>. 5
- [60] Christopher J. C. H. Watkins and Peter Dayan. Q-learning. *Mach. Learn.*, 8(3), 1992. 2
- [61] 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. *CoRR*, abs/2306.09341, 2023. 1, 5, 8, 13, 23

- [62] Xiaoshi Wu, Yiming Hao, Manyuan Zhang, Keqiang Sun, Zhaoyang Huang, Guanglu Song, Yu Liu, and Hongsheng Li. Deep reward supervisions for tuning text-to-image diffusion models. In *Proc. ECCV*, 2024. [2](#)
- [63] Zhengpeng Xie, Qiang Zhang, Fan Yang, Marco Hutter, and Renjing Xu. Simple policy optimization. In *Proc. ICML*, 2025. [5](#), [9](#), [20](#), [23](#)
- [64] 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 *Proc. NeurIPS*, 2023. [2](#)
- [65] Shuchen Xue, Chongjian Ge, Shilong Zhang, Yichen Li, and Zhi-Ming Ma. Advantage weighted matching: Aligning RL with pretraining in diffusion models. *CoRR*, abs/2509.25050, 2025. [2](#)
- [66] Zeyue Xue, Jie Wu, Yu Gao, Fangyuan Kong, Lingting Zhu, Mengzhao Chen, Zhiheng Liu, Wei Liu, Qiushan Guo, Weilin Huang, and Ping Luo. DanceGRPO: Unleashing GRPO on visual generation. *CoRR*, abs/2505.07818, 2025. [1](#), [2](#), [3](#), [21](#)
- [67] Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. STaR: Bootstrapping reasoning with reasoning. In *Proc. NeurIPS*, 2022. [2](#)
- [68] Dinghuai Zhang, Yizhe Zhang, Jiatao Gu, Ruixiang Zhang, Joshua M. Susskind, Navdeep Jaitly, and Shuangfei Zhai. Improving GFlowNets for text-to-image diffusion alignment. *TMLR*, 2025. [2](#)

A. Additional results

Training progression and artifacts Fig. 10 shows examples of RL post-training progression (without classifier-free guidance) for our method and Flow-GRPO [34]. Our method converges much faster and its subjective image quality peaks around 200 epochs. Training much longer than that results in loss of diversity and monotonically progressing over-saturation of colors.

Meanwhile, Flow-GRPO slowly improves for approximately 500 epochs, and after that, its characteristic artifacts start to appear. These include grid-like noisy patterns that fade in and out during training, as well as random changes to image styles. Fig. 12 shows vertical stripes at epoch 600 and particularly strong horizontal stripes at epoch 830. The artifacts are tied to the particular epochs and show up in almost all prompts. The presence of artifacts can vary quickly. For example, the artifacts in epoch 830 have again disappeared just 5 epochs later at epoch 835. Similarly, the style of the images can change quickly. Epochs 670 and 1000 prefer simplistic cartoony style across most prompts.

Comparison to classifier-free guidance (CFG) Fig. 11 shows RL post-training progressions with CFG enabled for our method and Flow-GRPO. The starting points are of vastly higher quality than without CFG (Fig. 10), and RL post-training has only a limited effect, mainly increasing the amount of background detail slightly.

Fig. 13 shows a direct comparison between our RL post-training and CFG. RL post-training leads to significantly more detailed images at the same level of diversity.

Combining rewards We find VLM-based rewards to be a promising way to guide RL post-training. They offer considerable freedom in designing the visual look without collecting additional data, and are likely to keep improving in the future as more capable off-the-shelf VLMs become available.

Here, we experiment with various combinations of VLM-based rewards and the more targeted PickScore reward [27] that is designed to mimic human preferences. Fig. 14 shows the effect of different post-training reward functions on the visual look using a number of prompts. Details of the reward functions are as follows:

- a) PickScore as the only reward function
- b) $\text{PickScore} + \text{VLM alignment} \times 0.1$
“Does this image match the caption “...”? Answer Yes or No.”
- c) VLM alignment as the only reward function
- d) $\text{VLM alignment} + \text{VLM quality} \times 0.4$
“Is this image of professional quality? Answer with Yes or No.”
- e) $\text{VLM alignment} + \text{VLM photorealism} \times 0.4$
“Does this image look photorealistic? Answer Yes or No.”

Using PickScore reward alone produces a somewhat muted and recognizable AI-generated look, which is improved with the addition of the VLM-based prompt alignment goal. However, VLM prompt alignment alone does not define a consistent style and suffers from obvious quality issues. Adding a goal that asks for professional quality often leads to non-photorealistic styles, whereas asking for photorealism leads to mostly realistic lighting and a consistent visual look.

We note that computing the VLM reward components as separate additive terms seemed to work much better in our tests than combining multiple objectives in a single prompt. It appeared that in the latter case, the objective that was easiest to quantify tended to dominate the overall assessment of the image, so that the less clear-cut objectives became largely neglected during post-training.

Additional metrics Fig. 15 provides a range of benchmark results for our method and Flow-GRPO, using models post-trained using the combined reward. A number of observations can be made.

The constituent parts of the combined reward (PickScore and VLM alignment) are both successfully optimized, rising quickly for us and slower for Flow-GRPO. For our method PickScore starts to drop in the over-training regime where colors are exceedingly saturated.

HPSv2 [61] is an alternative human-preference model, and it tells a very similar story compared to PickScore. CLIP models [47] do not directly represent human preferences, but also tell a substantially similar story, albeit with a twist that CFG is more strongly preferred.


OneIG-Bench alignment [7] is highly correlated with our VLM alignment, and OneIG-Bench diversity closely matches the widely used DreamSim diversity [17]. OneIG-Bench text and style are only tangentially related to the combined reward, and thus weakly affected by the optimization. OneIG-Bench reasoning is another benchmark that seems to prefer CFG.

See Appendix B.4 for further details on the evaluation protocol.

KL regularization Fig. 16 illustrates that our method is compatible with KL regularization. The primary goal of KL is to preserve more of the diversity, and we can see that it clearly achieves this. The results depend strongly on the regularization strength, which acts as a tradeoff between the retained diversity and prompt/human preference alignment. The tradeoff is qualitatively similar between our method and Flow-GRPO, although our method would need a lower regularization strength to reach its optimal tradeoff.

We implement KL regularization as an L2 distance penalty between the velocity predicted by the base model and model being finetuned with RL. This matches the implementation in prior works [15, 34].

Our method without CFG



Flow-GRPO without CFG




Figure 10. Qualitative comparison without classifier-free guidance, using prompt “A cat-dragon hybrid. Photograph.”

Our method with CFG

Seed A						
						
						
	Epoch: 0 Reward: 27.61 Alignment: 78.88 Diversity: 26.61	Epoch: 50 Reward: 28.52 Alignment: 81.70 Diversity: 22.88	Epoch: 100 Reward: 29.32 Alignment: 83.57 Diversity: 19.62	Epoch: 200 Reward: 29.91 Alignment: 85.22 Diversity: 16.89	Epoch: 500 Reward: 30.70 Alignment: 86.50 Diversity: 13.05	Epoch: 1000 Reward: 30.46 Alignment: 84.85 Diversity: 12.04

Flow-GRPO with CFG




Seed A						
						
						
	Epoch: 0 Reward: 27.68 Alignment: 78.96 Diversity: 26.14	Epoch: 50 Reward: 27.80 Alignment: 78.79 Diversity: 25.35	Epoch: 100 Reward: 28.05 Alignment: 79.42 Diversity: 23.89	Epoch: 200 Reward: 28.35 Alignment: 80.50 Diversity: 21.76	Epoch: 500 Reward: 29.02 Alignment: 81.72 Diversity: 21.29	Epoch: 1000 Reward: 29.75 Alignment: 83.14 Diversity: 17.03

Figure 11. Qualitative comparison with classifier-free guidance, using prompt “A cat-dragon hybrid. Photograph.”




Figure 12. Flow-GRPO results on selected prompts and epochs. Combined reward, no CFG, no KL. (a) “A girl with pigtaails is holding a giant sunflower.” (b) “Lunch in Bavaria - oil painting” (c) “A cat dressed as a wizard in broad daylight.” (d) “Budapest as a beautiful flowerpunk city, flowerpunk, hyper realistic, high quality, 8k” (e) “Kids race their bikes down the hill as their friends cheer from the sidelines, and a kite flutters in the breeze above them.” (f) “A quiet suburban cul-de-sac, where children play in its enclosed street.”




Figure 13. Qualitative comparison between our RL post-training and classifier-free guidance (CFG). On each row, the OneIG-Bench diversity matches between the two methods.




Figure 14. Effect of different reward functions on the visual look. Each column was post-trained for 200 epochs using our method without KL regularization or CFG.




Figure 15. A wide range of benchmarks for four models post-trained using the combined reward.




Figure 16. Effect of KL regularization on prompt alignment, human preferences, and diversity. These results are without CFG and with KL ratio $\beta = 0.04$, which is the default in Flow-GRPO.

Algorithm 2 Our training algorithm

```

procedure SAMPLEROLLOUTS( $\theta$ )
   $C \leftarrow \text{SamplePrompts}(n_{\text{prompts}})$  ▷ Draw a set of prompts
   $S \leftarrow \emptyset$  ▷ Initialize set of samples
  for each  $c \in C$  do ▷ Loop over prompts
    sample  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  ▷ Sample initial noise
     $\gamma_{0:T} \leftarrow \text{SampleStochSched}(T)$  ▷ Draw random
     $\hat{\gamma}_{0:T} \leftarrow \text{SampleStochSched}(T)$  stochasticity schedules
     $\mathbf{x}_{0:T}, \mathbf{v}_{0:T-1} \leftarrow \text{StochasticFlowSampler}(\theta, \epsilon, c, \gamma_{0:T})$ 
     $\hat{\mathbf{x}}_{0:T}, \hat{\mathbf{v}}_{0:T-1} \leftarrow \text{StochasticFlowSampler}(\theta, \epsilon, c, \hat{\gamma}_{0:T})$ 
     $\Delta R \leftarrow R(\hat{\mathbf{x}}_T, c) - R(\mathbf{x}_T, c)$  ▷ Reward difference
     $\Delta \mathbf{x} \leftarrow \hat{\mathbf{x}}_T - \mathbf{x}_T$  ▷ Image difference
     $\Delta \mathbf{x} \leftarrow \Delta \mathbf{x} / (\|\Delta \mathbf{x}\|_{\text{RMS}}^2 + 10^{-6})$  ▷ Normalize  $\Delta \mathbf{x}$ 
    for each  $t \in \{0, \dots, T-1\}$  do
       $S \leftarrow S \cup (\mathbf{x}_t, \mathbf{v}_t, t, c, \Delta R, \Delta \mathbf{x})$  ▷ Save all samples
       $S \leftarrow S \cup (\hat{\mathbf{x}}_t, \hat{\mathbf{v}}_t, t, c, \Delta R, \Delta \mathbf{x})$  from both trajectories
  return  $S$ 

procedure TRainePOCH( $\theta$ )
   $\text{samples} \leftarrow \text{SampleRollouts}(\theta)$  ▷ Sample rollouts
   $\text{batches} \leftarrow \text{RandomSplit}(\text{samples}, n_{\text{batches}})$  ▷ Split into batches
  for each  $\text{batch} \in \text{batches}$  do ▷ Loop over batches
     $g \leftarrow \mathbf{0}$  ▷ Gradient accumulator
    for each  $\text{sample} \in \text{batch}$  do
       $(\mathbf{x}, \mathbf{v}_{\text{ref}}, t, c, \Delta R, \Delta \mathbf{x}) \leftarrow \text{sample}$ 
       $\mathbf{v}_{\text{cur}} \leftarrow \mathbf{v}_{\theta}(\mathbf{x}; t, c)$ 
       $L \leftarrow \text{SPO}(\mathbf{v}_{\text{cur}}, \mathbf{v}_{\text{ref}}, \Delta R, \Delta \mathbf{x})$ 
       $g \leftarrow g + \nabla_{\theta} L$  ▷ Accumulate gradients
     $\theta \leftarrow \text{AdamWUpdate}(\theta, g)$  ▷ Apply gradients
  return  $\theta$ 

procedure SPO( $\mathbf{v}_{\text{cur}}, \mathbf{v}_{\text{ref}}, \Delta R, \Delta \mathbf{x}$ ) ▷ Adapted from [63]
   $\mathbf{v}_{\text{target}} \leftarrow \mathbf{v}_{\text{ref}} - \Delta \mathbf{x}$ 
   $\rho \leftarrow \exp(\|\mathbf{v}_{\text{target}} - \mathbf{v}_{\text{ref}}\|_2^2 - \|\mathbf{v}_{\text{target}} - \mathbf{v}_{\text{cur}}\|_2^2)$ 
   $L \leftarrow -\Delta R \cdot \rho - (\|\Delta R\|/(2\epsilon)) \cdot (\rho - 1)^2$ 
  return  $L$ 

```

B. Implementation details

In the following, we go into further detail on our stochastic sampler and training algorithms. In the full training pseudocode, presented in Algorithm 2, the stochastic sampler (Algorithm 1) is called twice in the rollout sampling loop to construct the pair of trajectories from a common initial noise.

Details of the stochastic sampler are discussed in Section B.1. Calls to the “SampleStochSched” procedure in Algorithm 2 correspond to drawing random stochasticity schedules for the trajectories. These are discussed in Section B.2, followed by further notes on implementation details and evaluation in Sections B.3–B.5.

B.1. Stochastic sampling

Traditional Euler–Maruyama SDE solvers (e.g., [56]) come with drawbacks when applied to denoising diffusion. They effectively alternate between (1) removing noise by an ODE-like diffusion solver step, and (2) adding a fraction of the removed noise back as freshly drawn noise. However, the exact amount of noise removed and added in these steps is slightly off-sync due to discretization errors, leading to a discrepancy that compounds over multiple sampling steps (see Karras et al. [24] for a detailed discussion). Furthermore, the SDE-centric approach binds the noise addition rate to the underlying diffusion time and other incidental parameterization details, rather than to the steps themselves. Some steps might experience very large noise injections that are numerically problematic as they extend the length of the associated ODE-like step, while others get vanishingly small noises.

We build our stochastic flow matching sampler based on the EDM stochastic sampler [24]. This sampler departs from strict adherence to an SDE and explicitly implements the alternation between the two sub-steps as noise addition and an ODE solver step. The key insight is to interpret the noise addition sub-step as a genuine jump to a higher value of the time parameter, maintaining the correct marginal distribution. This ensures that the correct information about the noise level, signal scale, and the time parameter are passed to the network at the ODE solver sub-step. Furthermore, the noise injections are scaled in proportion to the existing noise level in the sample.

Specifically, the EDM sampler works by overshooting

the ODE step to a lower noise level than the deterministic sampler step schedule would indicate (roughly, by a fraction γ_i of the noise level, where γ_i is a per-step hyperparameter), and then adding just the right amount of fresh noise to hit the target.

Adaptation to flow matching When using unit Gaussian noise as the latent distribution, flow matching can be seen as a special case of a diffusion model. The usual derivation (e.g., [35]) focuses on the deterministic ODE sampling procedure. Extensions to an SDE have since been constructed (e.g., [34, 66]) but they suffer from the practical shortcoming discussed above. We shall thus apply the idea behind EDM sampler directly to flow matching. The same procedure can be used to derive an EDM-style stochastic sampler for any diffusion schedule (e.g., DDPM [21], DDIM [53], VP-SDE and VE-SDE [56], etc.) expressed as $\sigma(t)$ and $s(t)$ by following the same steps.

EDM parameterization We first note that flow matching can be expressed in the general EDM parametrization (see Table 1 in [24]) as having:

- Scale schedule $s(t) = 1 - t$, corresponding to the signal linearly fading out over the unit time interval from $[0, 1]$.
- Noise schedule $\sigma(t) = t/(1 - t)$, corresponding to the “noise-to-signal-ratio” rising from 0 to ∞ over the same interval.

As such, the apparent standard deviation of the noise in the intermediate noisy images, after scaling by $s(t)$, increases linearly with t , as expected for flow matching.

Let us construct an EDM-like stochastic step from time t_i to t_{i+1} so that we replicate the meaning of the hyperparameter γ_i from EDM. For practical implementation reasons related to the `diffusers` library abstractions, we take the noise reduction step first and add noise second, while EDM takes these sub-steps in the opposite order. We also use a 1st-order Euler ODE solver instead of the 2nd-order Heun scheme [4] originally used in EDM.

Noise reduction sub-step The deterministic ODE sampler would simply take an Euler step from time t_i and t_{i+1} and proceed to the next step. In the stochastic version, we overshoot the target time t_{i+1} to a less noisy time $\tilde{t}_{i+1}$. For this, we must calculate the corrected target noise level.

The EDM sub-step calls for dividing the target noise level by $1 + \gamma_i$, where $\gamma_i \geq 0$. Substituting this into the schedule, this gives the overshoot target noise level

$$\tilde{\sigma}_{i+1} = \frac{\sigma(t_{i+1})}{1 + \gamma_i} = \frac{t_{i+1}}{\gamma_i - \gamma_i t_{i+1} - t_{i+1} + 1}. \quad (3)$$

Then, to find the corresponding ODE time, we invert the noise schedule formula, giving $t(\sigma) = \sigma/(\sigma + 1)$. Substituting $\tilde{\sigma}_{i+1}$, the overshoot time becomes

$$\tilde{t}_{i+1} = \frac{t_{i+1}}{1 - \gamma_i t_{i+1} + \gamma_i}. \quad (4)$$

The first half of Algorithm 1 uses this calculation to take an Euler step for the current image $\mathbf{x}_i$ from time t_i to $\tilde{t}_{i+1}$. Following the ODE velocity for this interval reduces the noise level and scales up the signal accordingly, yielding the intermediate image $\tilde{\mathbf{x}}_{i+1}$.

Noise addition sub-step Next, we determine the amount of noise to add in order to reach the noise level corresponding to the endpoint of the solver step, i.e., $\sigma(\tilde{t}_{i+1})$. Because of the non-uniform scale schedule employed by flow matching, we also need to manually re-scale the resulting mixture of signal and noise to match the expected scale at t_{i+1} . Given the intermediate image $\tilde{\mathbf{x}}_{i+1}$ and freshly drawn unit Gaussian noise $\epsilon_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, the target image is a linear mixture of these with weights that we shall now derive.

The non-uniform scale schedule employed by flow matching makes direct calculation cumbersome. Let us temporarily remove it by dividing out the scale of the current sample as $\tilde{\mathbf{x}}_{i+1}/s(\tilde{t}_{i+1})$. Then, the noise level $\sigma(\tilde{t}_{i+1})$ directly indicates the standard deviation of the noise present in this image. To reach the noise level $\sigma(t_{i+1})$, we must add enough noise to cover the difference in the variance (square) of the two noise levels, as variance between independent noises adds linearly.

The gap in variance to cover is $\sigma^2(t_{i+1}) - \sigma^2(\tilde{t}_{i+1})$, and the standard deviation of the noise needed to do so is the square root of this quantity. Thus, the unscaled sample with the added noise is $\tilde{\mathbf{x}}_{i+1}/s(\tilde{t}_{i+1}) + \sqrt{\sigma^2(t_{i+1}) - \sigma^2(\tilde{t}_{i+1})}\epsilon_i$. Finally, we reinstate the flow matching scaling at the target noise level by multiplying this mixture by $s(t_{i+1})$. Substituting everything, the formula for the noise addition and scaling comes out as

$$\mathbf{x}_{i+1} = \frac{\tilde{\mathbf{x}}_{i+1} + \tilde{t}_{i+1} \sqrt{\gamma_i^2 + 2\gamma_i} \cdot \epsilon_i}{\gamma_i \tilde{t}_{i+1} + 1}, \quad (5)$$




which corresponds to the second half of our stochastic sampling loop (Algorithm 1).

B.2. Stochasticity schedules

The main role of stochasticity in our method is to generate limited perturbations to an image during the pairwise sampling process. Given the flexibility of specifying the per-step stochasticity strength parameter γ_i (i.e., the stochasticity schedule) and the possibility of randomizing this choice between rollouts, there are several ways to achieve this. As it is not obvious which approach is preferable, we experimented with three families of randomized stochasticity schedules:

- **Uniform:** Specify a uniform, fixed value of γ_i used at all time steps for all rollouts. We use the value 0.0025 in all our experiments.
- **Interval:** For each rollout, pick a random center and width of a smooth interval in t . This results in both weak and strong perturbations, depending on the randomly chosen

	Schedule		OneIG alignment	
	Stochasticity	Weighting	200 epochs	Highest
U	Uniform	–	82.13	84.25
UL	Uniform	Low	80.47	83.67
UH	Uniform	High	81.45	84.37
I	Interval	–	82.29	83.50
IL	Interval	Low	76.41	80.09
IH	Interval	High	77.47	82.51
P	Prior	–	70.79	80.51
PL	Prior	Low	75.01	79.19
PH	Prior	High	79.58	82.99



(a) Exploration of different stochasticity schedules

(b) Uniform schedule

(c) Interval schedule

(d) Prior schedule

Figure 17. (a) Results for our three stochasticity schedules using the combined reward, with optional gradient weighting. We report the OneIG-Bench alignment at 200 epochs as well as the highest achieved score. Our main method corresponds to the first row. (b–d) Convergence of OneIG-Bench alignment for each schedule, annotated with dots that correspond to the numbers shown in the table.

interval width. The perturbations affect the image details relevant to the noise levels within the chosen interval.

- **Prior:** Make a random perturbation of randomized magnitude at the initial noise level, i.e., at the prior distribution, and run the rest of the sampling trajectory deterministically.

We also experimented with applying the training gradients using non-uniform weights per time step, as opposed to weighting them equally at all time steps. The hypothesis was that it might be beneficial to focus the updates primarily to some specific range of steps. We tried three variants for each schedule: uniform weighting, focusing low noise levels, and focusing high noise levels. Specifically, focusing low noise levels assigns a weight for each t according to density function $\mathcal{LN}(t; -0.3, 1.0)$ and focusing high noise levels according to $\mathcal{LN}(t; 0.3, 1.0)$, where $\mathcal{LN}(x; \mu, \sigma)$ is the logit-normal distribution:

$$\mathcal{LN}(x; \mu, \sigma) := \frac{1}{\sqrt{2\pi}\sigma x(1-x)} \exp\left[-\frac{(\text{logit}(x) - \mu)^2}{2\sigma^2}\right], \quad (6)$$

where

$$\text{logit}(x) := \log \frac{x}{1-x} \quad (7)$$

is the inverse function of a sigmoid. After evaluating the gradient weights for each t , they are normalized to sum to 1.

Fig. 17 shows prompt alignment and diversity measurements for the three stochasticity schedules with the three gradient weighting options. The uniform schedule is largely unaffected by gradient weighting and is the best option overall. The other two schedules are quite sensitive to gradient weighting and yield inferior results.

Fig. 18 shows a small set of representative visual results for selected configurations. Our observations below about the subjective image quality for each configuration are based on the results at large, not only to the images included in the figure. Our baseline method, the uniform schedule without

gradient weighting (column U), results in good image quality, variation, and tonal balance. Enabling gradient weighting (column UH) causes the images to be consistently too dark, with somewhat reduced contrast.

Interval schedule The interval schedule is configured by hyperparameters μ_{center} and σ_{center} for specifying the randomization of the interval center position, σ_{int} for interval width, and w_{int} for overall strength. After the interval center has been determined, γ_i are assigned using the logit-normal distribution density (Eq. 6). Denoting the vector of timesteps $[t_0, t_1, \dots, t_{T-1}]$ as $\mathbf{t}$ and the vector of outputs $[\gamma_0, \gamma_1, \dots, \gamma_{T-1}]$ as $\boldsymbol{\gamma}$, the per-rollout stochasticity schedule is drawn as follows.

Algorithm 3 Interval stochasticity schedule

sample $c \sim \mathcal{N}(\mu_{\text{center}}, \sigma_{\text{center}})$	▷ Random interval center
$\boldsymbol{\gamma} \leftarrow \mathcal{LN}(\mathbf{t}; c, \sigma_{\text{int}})$	▷ Evaluate LN density
$\boldsymbol{\gamma} \leftarrow \boldsymbol{\gamma} / \text{sum}(\boldsymbol{\gamma})$	▷ Normalize
$\boldsymbol{\gamma} \leftarrow \exp(w_{\text{int}} \cdot \boldsymbol{\gamma}) - 1$	▷ Apply constant weight
return $\boldsymbol{\gamma}$	

The normalization ensures that γ_i sum to 1 prior to applying the overall strength regardless of, e.g., spacing of the time steps within the interval. The exponentiation and subtraction of 1 on the last line contribute a minor correction that equalizes the total impact of stochastic randomization when it is distributed between several steps.

In our experiments we use hyperparameter values $\mu_{\text{center}} = 1.3$, $\sigma_{\text{center}} = 1.5$, $\sigma_{\text{int}} = 0.25$, and $w_{\text{int}} = 3$. In a separate experiment, we tried matching the gradient weights to the drawn stochasticity interval, but found this to perform significantly worse.

The interval schedule without gradient weighting (Fig. 18, column I) yields good subjective quality and variation, and excellent contrast. It produces cartoony style more often than the uniform schedule. The high quality is also confirmed by the high average alignment scores in Fig. 17a, row I. With gradient weighting (Fig. 18, column IH) the images become

clearly over-exposed with simplified details, overdoing the cartoony look.

Prior schedule The prior schedule is configured via hyperparameters $\mu_{\log w}$ and $\sigma_{\log w}$ that specify the distribution of the strength of stochasticity in each rollout. Using log-normal distribution here ensures that the final applied strength is always positive.

Algorithm 4 Prior stochasticity schedule

sample $W \sim \mathcal{N}(\mu_{\log w}, \sigma_{\log w})$	▷ Sample log of weight
$w \leftarrow \exp W$	▷ Random positive weight
$\gamma \leftarrow 0$	▷ Zero stochasticity
$\gamma_0 \leftarrow \exp(w) - 1$	▷ .. except at first time step
return γ	

In our experiments, we use values $\mu_{\log w} = \log 0.1$ and $\sigma_{\log w} = 1.0$.

Prior schedule without gradient weighting (Fig. 18, column P) produces low-detail images with water color look to them. The contrast and exposure are excessively high. This is also confirmed by the low average alignment scores in Fig. 17a, row P. Enabling gradient weighting (column PH) improves the results significantly, yielding a model that has decent image quality and variation. That said, the images are slightly under-exposed, and there is a tendency to exaggerate background details as well.

B.3. Hyperparameters and practical considerations

We use the hyperparameters specified in Table 1 across our main experiments and ablations when running variants of our method. As explained in Section 4.3, we effectively process four training batches of 8640 samples/batch per epoch. In practice, we do this by running 40 sub-batches with 216 samples each to conserve memory, accumulating gradients until the batch is complete, after which the gradients are used to update the model weights.

Ratio clipping An important detail is how we implement PPO-style ratio clipping. We cannot use tractable likelihoods in our algorithm, as these are generally inaccessible in diffusion and flow models, so we compute a proxy ratio for clipping. Our gradient for the velocity takes the form $\Delta R \Delta \bar{\mathbf{x}}$, where $\Delta \bar{\mathbf{x}} = \text{normalize}(\hat{\mathbf{x}}_T - \mathbf{x}_T)$ is the image difference and $\Delta R = R(\hat{\mathbf{x}}_T) - R(\mathbf{x}_T)$ is the reward difference. This is analogous to an advantage estimate (ΔR) times the gradient of a log likelihood ($\Delta \bar{\mathbf{x}}$) in a policy gradient framework.

We compute a target velocity from each velocity prediction saved during rollouts, $\mathbf{v}_{\text{target}} \leftarrow \mathbf{v}_{\text{ref}} - \Delta \bar{\mathbf{x}}$, and compose an L2 loss between this target velocity and the model’s current velocity prediction, which is analogous to the log-likelihood in Gaussian PPO or the conditional flow matching loss in FPO [39]. We compute our ratio as the exponentiated difference of this loss with the old policy parameters and the current policy parameters, i.e.,

Hyperparameter	Value
<i>Model Settings</i>	
LoRA rank	32
LoRA alpha	64
Number format	fp16 (+ scaling)
<i>Training Settings</i>	
Prompts/pairs per epoch	432
Sampling trajectories per epoch	864
Denoising steps per trajectory	40
Denoising steps per epoch	$40 \times 864 = 34,560$
Training batch size	216
Training batches per epoch	160
Gradient accumulation	40 batches
Model weight updates per epoch	4
Clip range	$3 \cdot 10^{-2}$
Learning rate	$3 \cdot 10^{-5}$
Weight decay	$1 \cdot 10^{-4}$
Gradient clip norm	1
Ratio clipping	SPO [63]
Advantage/reward clipping	None
KL early stopping	None
<i>Evaluation Settings</i>	
ODE step count	40
ODE step type	Euler, deterministic
ODE step sizes	Logit-normal (SD3.5 [14])

Table 1. RL training hyperparameters for our method.

$\exp(\|\mathbf{v}_{\text{target}} - \mathbf{v}_{\text{ref}}\|_2^2 - \|\mathbf{v}_{\text{target}} - \mathbf{v}_{\text{cur}}\|_2^2)$. This is a ratio in the same form as the PPO/SPO likelihood ratio, so we can apply clipping in an identical manner. The “SPO” procedure in Algorithm 2 provides a pseudocode implementation.

B.4. Evaluation protocol

All metrics reported in this paper were calculated *post hoc* based on LoRA checkpoints exported during the corresponding RL training runs. Specifically, our results never include reward values seen during training, as these could be biased by, e.g., the perturbations to the sampling trajectories.

For the reward and its individual components (e.g., Fig. 15{a,b,c}), we sample 4,096 random prompts from the Pick-a-Pic [27] training set, which we process the same way as the official Flow-GRPO implementation to yield a total of 25,432 prompts. To avoid selection bias, we re-randomize the choice of prompts and the per-image random seeds on a per-checkpoint basis. For consistency with our VLM rewards, we scale the raw value of PickScore by 100.

For HPSv2, CLIP-H/14, CLIP-L/14, and DreamSim diversity (Fig. 15{d,e,f,i}), we run the official models¹²³⁴ with all 3,200 prompts from the HPDv2 dataset [61], generating 4 images for each prompt using different random seeds to improve accuracy. In line with previous work, we define DreamSim diversity of a pair of images as the sum of

¹<https://huggingface.co/adams-story/HPSv2-hf>

²<https://huggingface.co/laion/CLIP-ViT-H-14-laion2B-s32B-b79K>

³<https://huggingface.co/openai/clip-vit-large-patch14>

⁴<https://github.com/ssundaram21/dreamsim>




Figure 18. Effect of different stochasticity schedules (from Fig. 17) on the visual look of different prompts. Each column was post-trained for 200 epochs using our method with the combined reward without CFG.

Item	H200 benchmark		Ours, 40 steps		Flow-GRPO, 40 steps		Ours, 10 steps		Flow-GRPO, 10 steps	
	ms/eval	batch	eval/epoch	s/epoch	eval/epoch	s/epoch	eval/epoch	s/epoch	eval/epoch	s/epoch
Denoiser fwd	13.2	32	69120	912.4	68256	901.0	17280	228.1	16416	216.7
Denoiser bwd	17.9	32	34560	618.6	33696	603.2	8640	154.7	7776	139.2
VAE decoder	11.0	32	864	9.5	864	9.5	864	9.5	864	9.5
PickScore	18.8	32	864	16.2	864	16.2	864	16.2	864	16.2
Qwen-7B	35.3	26	864	30.5	864	30.5	864	30.5	864	30.5
Estimated total				1587.2		1560.4		439.0		412.1
Measured total				1625.5		2736.7		479.4		768.4
Overhead				+2.4%		+75.4%		+9.2%		+86.5%

Table 2. Details of the wall-clock performance results presented in Section 5.5.

squares between their normalized DreamSim embeddings, and average the results over every unique pair of images that were generated using the same prompt.

For OneIG-Bench (e.g., Fig. 15{g,h,j,k,l}), we use the official implementation⁵ with all 1,120 prompts and 4 images per prompt. In terms of reporting, we multiply the raw metric values by 100 for convenience.

B.5. Wall-clock performance

In Section 5.5, we compare the wall-clock performance of our method to Flow-GRPO. As our goal is to compare the two methods, not their implementations, we disregard implementation-related overheads and report ideal performance estimates assuming zero overhead. In practice, our implementation matches these ideal estimates closely, whereas the official implementation of Flow-GRPO is considerably less efficient. Comparison based on true wall-clock time would thus be needlessly unfair to Flow-GRPO.

The details of our performance estimation procedure are shown in Table 2. We first benchmark each major computational component in isolation on NVIDIA H200 using representative inputs and reasonably sized batches. For example, running the denoiser (“Denoiser fwd”) for a batch of 32 noisy latent images, corresponding to 512×512 resolution in terms of RGB pixels, takes approximately 13.2 ms per image on average, and backpropagating the resulting gradients to the LoRA weights (“Denoiser bwd”) takes a further 17.9 ms per image. A similar measurement is carried out for the remaining components, including the evaluation of the reward functions.

We then instrument each implementation to report the number of times each underlying component is evaluated per epoch. Tallying everything together, we arrive at the numbers shown in the “Estimated total” row, which we use to rescale the x -axis for each curve in Fig. 9b.

For reference, we also measure the true wall-clock time per epoch taken by each method, averaged over the entire RL training run. Comparing these timings (“Measured total”) against the implementation-neutral estimates, we see that our implementation overhead, including logging and

checkpoint export, is less than 10% in both 40 and 10 step cases, whereas the official Flow-GRPO implementation has overheads exceeding 75%.

C. Theoretical analysis of our method

In this section, we develop a connection between our finite differences and an approximate gradient of a smoothed version of the reward.

C.1. Reward gradient

As noted in Section 3.1, we assume there is a reward function $R(\mathbf{x})$ that maps images to scalar reward values, with higher values representing more desirable images. Our goal is to fine-tune the pre-trained velocity function v_θ to maximize the expected reward over draws from the generative model

$$\arg \max_{\theta} \mathbb{E}_{\mathbf{c} \sim C, \mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} R(f_\theta(\mathbf{x}_0; \mathbf{c})), \quad (8)$$

where f_θ refers to the entire sampling process that utilizes the learned velocity function v_θ .

If R were differentiable, we could ascend this objective by taking gradient ascent steps with respect to the weights θ . The gradient is the standard backpropagation

$$\mathbf{J}_f^\theta(\mathbf{x}_0; \mathbf{c})^\top \nabla R(f_\theta(\mathbf{x}_0; \mathbf{c})), \quad (9)$$

where $\mathbf{J}_f^\theta$ denotes the Jacobian matrix of f_θ with respect to θ , and the gradient ∇ is taken with respect to an image at time t_T . As each Euler ODE step contributes to the shared weight gradient additively, this further decomposes to a sum of terms contributed by individual timesteps

$$\sum_{i=1}^T (t_i - t_{i-1}) \mathbf{J}_v^\theta(\mathbf{x}_{i-1}; t_{i-1}, \mathbf{c})^\top \mathbf{J}_{f_i}^\mathbf{x}(\mathbf{x}_i; \mathbf{c})^\top \nabla R(\mathbf{x}_T). \quad (10)$$

Here, $\mathbf{J}_v^\theta$ is the Jacobian matrix of v_θ with respect to θ , f_i denotes running the ODE solver starting from an intermediate image at step i to completion, and $\mathbf{J}_{f_i}^\mathbf{x}$ is the Jacobian of this process with respect to said intermediate input image. In other words, the gradient of the reward at the generated image is backpropagated to the step i through the ODE chain, yielding a “gradient image” of same format as the noisy image, and then further backpropagated through the step velocity into the weights.

⁵<https://github.com/OneIG-Bench/OneIG-Benchmark>

C.2. Ascending the reward by finite differences

We cannot generally expect reward functions to be differentiable, and even when they are, their raw derivatives can behave poorly, e.g., by pointing to effective reward hacking directions. Our method performs an approximate optimization step that implicitly smooths the reward function and does not require direct gradient information, drawing inspiration from finite differences.

In the following, we analyze a prototypical version of our method where some training-technical practicalities and efficiency improvements have been dropped. In this variant, we generate a pair of similar but slightly perturbed images and train the flow to point towards the higher-reward one:

1. Draw an initial random noise image $\mathbf{x}_0$, and generate the full sequence of intermediate noise images $\{\mathbf{x}_i\}_{i=1}^T$ by solving the ODE steps.
2. Choose a random timestep j in $[1, T]$ to train at.
3. Make a normal-distributed perturbation $\mathbf{x}'_j \sim \mathcal{N}(\mathbf{x}_j, \sigma_c^2 \mathbf{I})$ of chosen scale σ_c around the noisy image $\mathbf{x}_j$ at step j .
4. Solve the remainder of the ODE starting at the perturbed image to obtain a second generated image $\mathbf{x}'_T$.
5. Form an approximate gradient

$$\tilde{g} = \frac{R(\mathbf{x}'_T) - R(\mathbf{x}_T)}{\sigma_c^2} (\mathbf{x}'_T - \mathbf{x}_T). \quad (11)$$

6. Backpropagate through the velocity at step $j - 1$ into weights using the standard procedure (Eq. 10), but with the backpropagated reward gradient replaced by $\tilde{g}$.

C.3. Approximate gradient

Let us now show that this prototypical algorithm approximates a gradient direction similar to the gradient of the reward. To reduce notational clutter, we will drop references to θ and prompt embeddings $\mathbf{c}$. As we focus on step j , we also denote f_j , $\mathbf{x}_j$ and $\mathbf{J}_{f_j}^{\mathbf{x}}$ as simply f , $\mathbf{x}$, and $\mathbf{J}$.

Consider the expectation of our reward-weighted difference around a given $\mathbf{x}$ perturbed by a Gaussian noise of standard deviation σ , and taken through the ODE flow

$$\mathbb{E}_{\epsilon \sim \mathcal{N}(0, \sigma^2 \mathbf{I})} [R(f(\mathbf{x} + \epsilon)) - R(f(\mathbf{x}))] [f(\mathbf{x} + \epsilon) - f(\mathbf{x})]. \quad (12)$$

Assuming that the perturbation is small, we can approximate the mapping of the flow by its first order Taylor expansion $f(\mathbf{x}) + \mathbf{J}\epsilon$ for a small perturbation ϵ at $\mathbf{x}$:

$$\approx \mathbb{E}_{\epsilon \sim \mathcal{N}(0, \sigma^2 \mathbf{I})} [R(f(\mathbf{x}) + \mathbf{J}\epsilon) - R(f(\mathbf{x}))] \mathbf{J}\epsilon. \quad (13)$$

The second term in the bracket vanishes, as $R(f(\mathbf{x}))$ is a constant with respect to the random variable, and the expectation of a linearly transformed normal distribution is zero.⁶

⁶The role of this term can be seen as variance reduction when estimating the expectation stochastically, as it provides a baseline reward for the center of the perturbations.

We then make a change of variables to absorb the linear transformation and offset into the random variable:

$$= \mathbb{E}_{\epsilon' \sim \mathcal{N}(f(\mathbf{x}), \sigma^2 \mathbf{J}\mathbf{J}^\top)} R(\epsilon') (\epsilon' - f(\mathbf{x})). \quad (14)$$

This form is amenable to applying the multivariate Stein’s lemma, yielding

$$= \sigma^2 \mathbf{J}\mathbf{J}^\top \mathbb{E}_{\epsilon' \sim \mathcal{N}(f(\mathbf{x}), \sigma^2 \mathbf{J}\mathbf{J}^\top)} \nabla R(\epsilon') \quad (15)$$

$$= \sigma^2 \mathbf{J}\mathbf{J}^\top \nabla \mathbb{E}_{\epsilon' \sim \mathcal{N}(f(\mathbf{x}), \sigma^2 \mathbf{J}\mathbf{J}^\top)} R(\epsilon') \quad (16)$$

$$= \sigma^2 \mathbf{J}\mathbf{J}^\top \nabla \tilde{R}(f(\mathbf{x})). \quad (17)$$

The expectation can thus be read as the gradient of a smoothed reward function $\tilde{R}$, where its value is averaged over a Gaussian kernel of covariance $\sigma^2 \mathbf{J}\mathbf{J}^\top$. The smoothing turns the potentially discontinuous reward into a differentiable one. The smoothed gradient is further transformed by $\mathbf{J}\mathbf{J}^\top$, where we can interpret $\mathbf{J}^\top$ as backpropagating the reward gradient through the ODE to the time of the perturbation — thus approximating the gradient of the combined flow and reward in Eq. 10 — and the extra $\mathbf{J}$ as distorting this gradient estimate.

C.4. Characterizing the distortion

We first note that in practice the Jacobian of a diffusion flow mapping tends to be close to a positive definite symmetric matrix, where $(\mathbf{J}q)^\top q \geq 0$ for any vector q . Then, $\mathbf{J}$ will typically map the gradient to the same half-space, and it remains an ascent direction for the reward. The approximate positive definiteness follows from $\mathbf{J}$ being a composition of gently time-varying, genuinely positive definite Jacobians of the infinitesimal ODE steps. The mapping established by diffusion is also typically close to an optimal transport coupling, which would have an exactly positive definite Jacobian everywhere.

In practice, many editing-flavored diffusion methods implicitly rely on an argument of this nature; they make coarse edits to the intermediate noisy images and expect them to retain their approximate pixel position and color in the final generated image, but in a more “fleshed out”, fully generated form. In a sense, the dot product between the coarse edit and the difference it makes on the final image is empirically expected to be positive, hinting at the Jacobian being positive definite for practical purposes.

C.5. Discussion

Equations 12 to 17, taken in isolation, imply the claim in Section 4.1 that

$$\mathbb{E}[\nabla R(f(\mathbf{x}))^\top \mathbf{J} [\Delta R \Delta \mathbf{x}]] \geq 0, \quad (18)$$

where the expectation is taken over the random perturbations. In other words, the intermediate-time update $\Delta R \Delta \mathbf{x}$,

pushed through the remaining flow by $\mathbf{J}$, yields a direction of increasing reward.

By Eq. 17, the left hand side becomes

$$\sigma^2 \nabla R(f(\mathbf{x}))^\top \mathbf{J} \mathbf{J} \mathbf{J}^\top \nabla \tilde{R}(f(\mathbf{x})) \quad (19)$$

which is approximately a quadratic form evaluated with a positive definite matrix $\mathbf{J} \mathbf{J} \mathbf{J}^\top$, and as such, has a nonnegative value.

These results rely on assumptions discussed in Section C.2. Our practical algorithm modifies this idealized setting in various ways, motivated by efficiency and simplicity. The ablation with stochastic schedules (Section B.2) probes a specific difference: our method applies stochasticity throughout the trajectories instead of just a single step, and applies the update on *all* steps along the trajectories. In contrast, the **Interval** schedule models a situation similar to the idealized method, where the paths only deviate on a localized time interval. The empirical results suggest that a variety of schemes work in practice, and the more “ideal” scheme does not enjoy a clear benefit despite being more amenable for direct mathematical analysis.