

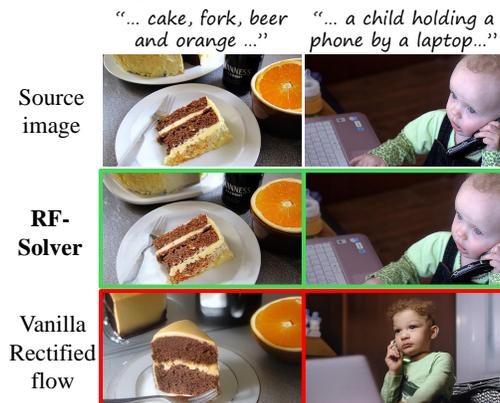
Taming Rectified Flow for Inversion and Editing

Jiangshan Wang^{1,2*}, Junfu Pu^{2*†}, Zhongang Qi^{2‡}, Jiayi Guo¹, Yue Ma³, Nisha Huang¹,
Yuxin Chen², Xiu Li^{1‡}, Ying Shan²

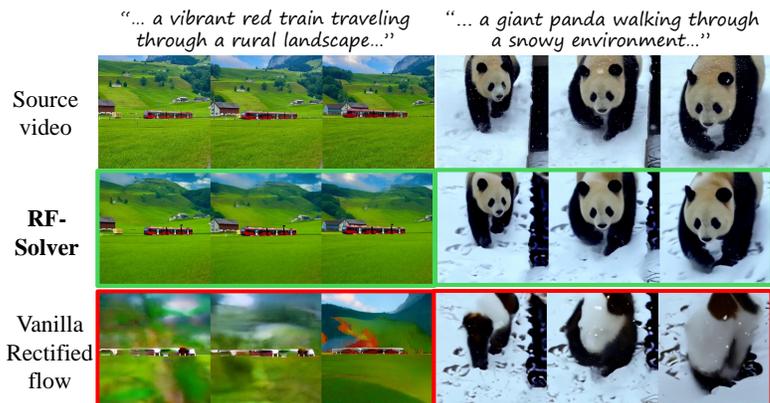
¹Tsinghua University ²ARC Lab, Tencent PCG ³HKUST

<https://rf-solver-edit.github.io>

Task 1: Image Inversion and Reconstruction



Task 2: Video Inversion and Reconstruction



Task 3: Image Editing (resolution: 1360 × 768)



Task 4: Video Editing (resolution: 1360 × 768; 96 frames)



Figure 1. **RF-Solver for downstream tasks in image and video.** We propose **RF-Solver** to solve the rectified flow ODE with reduced error, thereby enhancing both sampling quality and inversion-reconstruction accuracy for rectified-flow-based generative models [2, 3]. Furthermore, we propose **RF-Edit**, which utilizes the RF-Solver for image and video editing tasks. Our methods demonstrate impressive performance across generation, inversion, and editing tasks in both image and video modalities.

Abstract

Rectified-flow-based diffusion transformers like *FLUX* and *OpenSora* have demonstrated outstanding performance in the field of image and video generation. Despite their robust generative capabilities, these models often struggle with inversion inaccuracies, which could further limit their effectiveness in downstream tasks such as image and video editing. To address this issue, we propose *RF-Solver*, a

novel training-free sampler that effectively enhances inversion precision by mitigating the errors in the ODE-solving process of rectified flow. Specifically, we derive the exact formulation of the rectified flow ODE and apply the high-order Taylor expansion to estimate its nonlinear components, significantly enhancing the precision of ODE solutions at each timestep. Building upon *RF-Solver*, we further propose *RF-Edit*, a general feature-sharing-based framework for image and video editing. By incorporating self-attention features from the inversion process into the editing process, *RF-Edit* effectively preserves the structural information of the source image or video while achieving

*Equal contribution.

†Project lead.

‡Corresponding authors.

high-quality editing results. Our approach is compatible with any pre-trained rectified-flow-based models for image and video tasks, requiring no additional training or optimization. Extensive experiments across generation, inversion, and editing tasks in both image and video modalities demonstrate the superiority and versatility of our method. The source code is available at [this URL](#).

1. Introduction

Recent advancements of generation methods based on Rectified Flow (RF) [34, 41, 66] have demonstrated exceptional performance in synthesizing high-quality images and videos. Different from traditional approaches represented by Stable Diffusion [21, 55], these methods leverage the Diffusion Transformer [50, 62, 72, 77] architecture and implement a straight-line motion system to produce the desired data distribution. With these effective designs, FLUX [2] and OpenSora [3] have respectively emerged as one of the state-of-the-art (SOTA) methods in the field of Text-to-Image (T2I) and Text-to-Video (T2V) generation.

Despite the remarkable success in the fundamental T2I and T2V generation tasks, few studies have explored the performance of RF-based models on various downstream tasks such as inversion-reconstruction [18, 46, 59, 70] and editing [19, 44]. When directly applying the vanilla RF for inversion, we observe that it fails to faithfully reconstruct the image or video from the source. Examples are shown in Fig. 1 Task 1 and Task 2 (the third row). For image inversion, the positions of objects (*e.g.* the cake) and the appearance of individuals (*e.g.* the child) in the reconstructed image significantly diverge from the source image. The performance of video inversion is even worse, with noticeable distortions present in the reconstructed video. The inaccuracies of inversion and reconstruction would severely constrain the performance of RF models on other inversion-based downstream tasks such as image editing [14, 19, 28, 47, 63] and video editing [16, 32, 40, 58].

In this work, we investigate the aforementioned problem by delving into the inversion and reconstruction process of the RF. Specifically, we track the latent at each intermediate timestep during inversion and reconstruction, calculating the Mean Square Error (MSE) between them at corresponding timesteps. We observe that significant errors are introduced at each timestep throughout the whole reconstruction process, and their accumulation ultimately results in a considerably deviated output (the red curve in Fig. 2). Based on the definition and inference process of RF [39, 41], we identify that these errors stem from the Ordinary Differential Equation (ODE) solving process [22, 51, 65, 78]. Specifically, the essence of the inversion and generation process for RF is to derive the solution of RF ODE [41]. Since this ODE includes terms involving complex neural networks,

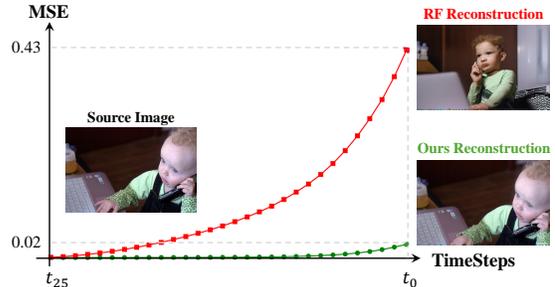


Figure 2. **Analysis of the inversion-reconstruction process.** Inversion takes the source image latent \tilde{Z}_{t_0} as the input and progressively add noise for N timesteps, obtaining $\tilde{Z}_{t_N} \in \mathcal{N}(0, I)$. \tilde{Z}_{t_N} is then denoised for N timesteps to obtain the reconstruction Z_{t_0} . During this process, we store the latent \tilde{Z}_{t_i} and Z_{t_i} at each timestep respectively in inversion and denoising processes. Then we calculate the Mean Squared Error (MSE) between them. The **red curve** represents the vanilla Rectified Flow inversion and the **green curve** represents RF-Solver inversion.

the solution can only be coarsely approximated by a sampler. However, the experiment in Fig. 2 indicates that the sampler adopted in existing models [2, 3] lacks sufficient precision for the inversion task, causing notable errors to accumulate at each timestep, finally leading to unsatisfactory reconstruction results.

Based on the analysis, we aim to improve inversion accuracy by introducing a more effective sampler, which is more general and fundamental than designing a specific inversion method. To this end, we propose **RF-Solver**. Specifically, we note that the exact formulation of the RF ODE solution can be directly derived using the variation of constants method. For the nonlinear component of this solution (*i.e.*, the integral of the neural network), we utilize Taylor expansion for estimation. By employing higher-order Taylor expansion, the ODE can be solved with reduced error, thereby enhancing the performance of RF models. RF-Solver is a generic sampler that can be seamlessly integrated into any rectified flow model without additional training or optimization. Experimental results demonstrate that RF-Solver not only significantly enhances the accuracy of inversion and reconstruction (the green curve in Fig. 2), but also improves performance on fundamental tasks such as T2I generation.

Building upon this, we propose **RF-Edit** to leverage RF-Solver in editing tasks. Real-world image and video editing require the model to make precise modifications to a source image/video while maintaining its overall structure unchanged, presenting greater challenges than reconstruction. In this scenario, it is inadequate to solely rely on the inverted noises as prior knowledge for editing, which could lead to edited results being excessively influenced by the target prompt, diverging significantly from the original source [19, 63]. Addressing this problem, RF-Edit stores the \mathcal{V} (value) feature in the self-attention layers at several timesteps during inversion. These features are used to re-

place the corresponding features in the denoising process. Practically, we design two specific sub-modules for RF-Edit, respectively leveraging the DiT structure of FLUX [2] and OpenSora [3] as the backbones for image and video editing. With the effective design of RF-Edit, it demonstrates superior performance in both image and video domains, outperforming various SOTA methods.

Our core contributions are summarized as follows:

- We propose RF-Solver, a training-free sampler that significantly reduces errors in the inversion and reconstruction processes of rectified-flow models.
- We present RF-Edit, which leverages RF-Solver for image and video editing. RF-Edit effectively preserves the structural integrity of the source image/video while delivering high-quality results.
- Extensive experiments across a range of tasks demonstrate the efficacy and versatility of our method.

2. Related Work

2.1. Inversion

Inversion maps the real visual data, *i.e.* image and video, to representations in noise space, which is the reverse process of generation. The representative method, DDIM inversion [59, 60], adds predicted noise recursively at each forward step. Many efforts [15, 26, 43, 45, 45, 46, 56, 64] have been made to mitigate the discretization error in DDIM inversion. Despite the effectiveness of inversion in diffusion models, the exploration of inversion in SOTA rectified flow models like FLUX and OpenSora is limited. RF-prior [75] uses the score distillation to invert the image while it requires many optimizing steps. More recently, [57] introduces an additional vector field conditioned on the source image to improve the inversion. However, the error from the original vector field of rectified flow still persists, which would limit the performance of such method on various downstream tasks. In contrast, we aim to directly mitigate the error from the original vector field in this work.

2.2. Image and Video Editing

Training-free methods for image and video editing [24, 61] have gained increasing popularity for their efficiency and effectiveness. Existing image editing methods focus on prompt refinement [54, 68], attention-sharing mechanism [6, 19, 49, 63], mask guidance [4, 11, 23, 36], and noise initialization [5, 76]. Video editing introduces additional complexities in maintaining temporal consistency, making it a more challenging task. Existing video editing methods focus on attention injection [40, 52, 69], motion guidance [10, 17, 67, 74], latent manipulation [9, 29, 73, 80], and canonical representation [8, 30, 35, 48]. To date, the editing performance of RF-based diffusion transformers has remained largely under-explored. Although [57] employs FLUX [2] for image editing, its performance is limited to

simple tasks such as stylization and face editing while often failing to effectively maintain the structural information of source images. Moreover, currently there is no research exploring the video editing capabilities of RF-based models.

3. Method

In this section, we present our method in detail. First, we introduce RF-Solver, which significantly enhances the precision of inversion and reconstruction. Subsequently, we present RF-Edit, an extension of RF-Solver designed to enable high-quality image and video editing.

3.1. Preliminaries

Rectified Flow (RF) [42] facilitates the transition between the real data distribution π_0 and Gaussian Noises distribution π_1 along a straight path. This is achieved by learning a forward-simulating system defined by the ODE: $d\mathbf{Z}_t = v(\mathbf{Z}_t, t)dt, t \in [0, 1]$, which maps $\mathbf{Z}_1 \in \pi_1$ to $\mathbf{Z}_0 \in \pi_0$.

In practice, the velocity field v is parameterized by a neural network v_θ . During training, given empirical observations of two distributions $\mathbf{X}_0 \sim \pi_0, \mathbf{X}_1 \sim \pi_1$ and $t \in [0, 1]$, the forward process (*i.e.* adding noise) of rectified flow is defined by a simple linear combination: $\mathbf{X}_t = t\mathbf{X}_1 + (1-t)\mathbf{X}_0$. The differential form of the equation is given by: $d\mathbf{X}_t = (\mathbf{X}_1 - \mathbf{X}_0)dt$. Consequently, the training process optimizes the network by solving the least squares regression problem, which fits the v_θ with $(\mathbf{X}_1 - \mathbf{X}_0)$:

$$\min_{\theta} \int_0^1 \mathbb{E} \left[\left\| (\mathbf{X}_1 - \mathbf{X}_0) - v_{\theta}(\mathbf{X}_t, t) \right\|^2 \right] dt. \quad (1)$$

In the sampling process, the ODE is discretized and solved using the Euler method. Specifically, the rectified flow model starts with a Gaussian noise sample $\mathbf{Z}_{t_N} \in \mathcal{N}(0, \mathbf{I})$. Given a series of N discrete timesteps $t = \{t_N, \dots, t_0\}$, the model iteratively predicts $v_{\theta}(\mathbf{Z}_{t_i}, t_i)$ for $i \in \{N, \dots, 1\}$ and then takes a step forward until generating the images \mathbf{Z}_{t_0} , with the following recurrence relation:

$$\mathbf{Z}_{t_{i-1}} = \mathbf{Z}_{t_i} + (t_{i-1} - t_i)v_{\theta}(\mathbf{Z}_{t_i}, t_i). \quad (2)$$

The RF model can generate high-quality images in much fewer timesteps compared to DDPM [21], owing to the nearly linear transition trajectory established during training. With these effective designs, RF model illustrates great potential in the field of T2I and T2V generation [2, 3].

3.2. RF-Solver

The vanilla RF sampler demonstrates strong performance in image and video generation. However, when applied to inversion and reconstruction tasks, we observe significant error accumulation at each timestep. This results in reconstructions that diverge notably from the original image (see Fig. 2). This severely limits the performance of RF models in various inversion-based downstream tasks [19, 65]. Delving into this problem, we identify that the errors stem from

the process of estimating the approximate solution for the rectified flow ODE [39, 66], which is formulated by Eq. (2) in existing methods [2, 3]. Consequently, obtaining more precise solutions for the ODE would effectively mitigate these errors, leading to improved reconstruction quality.

Based on this analysis, we start by carefully examining the differential form of the Rectified flow: $d\mathbf{Z}_t = \mathbf{v}_\theta(\mathbf{Z}_t, t)dt$. This ODE is discretized in the sampling process. Given the initial value \mathbf{Z}_{t_i} , the ODE can be exactly formulated using the *variant of constant* method:

$$\mathbf{Z}_{t_{i-1}} = \mathbf{Z}_{t_i} + \int_{t_i}^{t_{i-1}} \mathbf{v}_\theta(\mathbf{Z}_\tau, \tau) d\tau. \quad (3)$$

In the above formula, $\mathbf{v}_\theta(\mathbf{Z}_\tau, \tau)$ is the non-linear component parameterized by the complex neural network, which is difficult to approximate directly. As an alternative, we employ the Taylor expansion at t_i to approximate this term:

$$\mathbf{v}_\theta(\mathbf{Z}_\tau, \tau) = \sum_{k=0}^{n-1} \frac{(\tau - t_i)^k}{k!} \mathbf{v}_\theta^{(k)}(\mathbf{Z}_{t_i}, t_i) + \mathcal{O}((\tau - t_i)^n), \quad (4)$$

where $\mathbf{v}_\theta^{(k)}(\mathbf{Z}_{t_i}, t_i) = \frac{d^k \mathbf{v}_\theta(\mathbf{Z}_{t_i}, t_i)}{dt^k}$, denoting the k -order derivative of \mathbf{v}_θ and \mathcal{O} denotes higher-order infinitesimals. Substituting Eq. (4) into the integral term yields:

$$\int_{t_i}^{t_{i-1}} \mathbf{v}_\theta(\mathbf{Z}_\tau, \tau) d\tau = \sum_{k=0}^{n-1} \mathbf{v}_\theta^{(k)}(\mathbf{Z}_{t_i}, t_i) \int_{t_i}^{t_{i-1}} \frac{(\tau - t_i)^k}{k!} d\tau + \mathcal{O}((\tau - t_i)^n). \quad (5)$$

Through the above process, the network prediction term and its higher-order derivatives are separated from the integral. Then we notice that the remaining component in the integral can be computed analytically:

$$\int_{t_i}^{t_{i-1}} \frac{(\tau - t_i)^k}{k!} d\tau = \left[\frac{(\tau - t_i)^{k+1}}{(k+1)!} \right]_{t_i}^{t_{i-1}} = \frac{(t_{i-1} - t_i)^{k+1}}{(k+1)!}. \quad (6)$$

Substituting Eq. (6) and Eq. (5) into Eq. (3), we derive the n -th order solution of Rectified flow ODE:

$$\mathbf{Z}_{t_{i-1}} = \mathbf{Z}_{t_i} + \sum_{k=0}^{n-1} \frac{(t_{i-1} - t_i)^{k+1}}{(k+1)!} \mathbf{v}_\theta^{(k)}(\mathbf{Z}_{t_i}, t_i) + \mathcal{O}(h_i^{n+1}), \quad (7)$$

where $h_i := t_{i-1} - t_i$. Eq. (7) indicates that to estimate $\mathbf{Z}_{t_{i-1}}$, we need to obtain the k -th order derivatives $\{\mathbf{v}_\theta^{(k)}(\mathbf{Z}_{t_i}, t_i)\}$ for $k \in \{0, \dots, n-1\}$.

When $n = 1$, the formula reduces to the standard rectified flow (*i.e.*, Eq. (2)). In our experiments, we find that setting $n = 2$ effectively mitigates the errors, yielding:

$$\mathbf{Z}_{t_{i-1}} = \mathbf{Z}_{t_i} + (t_{i-1} - t_i) \mathbf{v}_\theta(\mathbf{Z}_{t_i}, t_i) + \frac{1}{2} (t_{i-1} - t_i)^2 \mathbf{v}_\theta^{(1)}(\mathbf{Z}_{t_i}, t_i). \quad (8)$$

Note that $\mathbf{v}_\theta^{(1)}$ in Eq. (8) is the first-order derivative of the network prediction term \mathbf{v}_θ , which cannot be analytically derived due to the complex architecture of the neural network. To estimate this term, we first obtain the network prediction $\hat{\mathbf{v}}_{t_i}$ at the timestep t_i , *i.e.*, $\hat{\mathbf{v}}_{t_i} = \mathbf{v}_\theta(\mathbf{Z}_{t_i}, t_i)$. Then we step forward a small timestep $\Delta t = \frac{1}{2}(t_{i-1} - t_i)$, and update the latents to obtain $\mathbf{Z}_{t_i + \Delta t} = \mathbf{Z}_{t_i} + \Delta t \cdot \hat{\mathbf{v}}_{t_i}$.

Subsequently, we calculate an additional prediction of the network at the timestep $t_i + \Delta t$, *i.e.*, $\hat{\mathbf{v}}_{t_i + \Delta t} = \mathbf{v}_\theta(\mathbf{Z}_{t_i + \Delta t}, t_i + \Delta t)$. With $\hat{\mathbf{v}}_{t_i}$ and $\hat{\mathbf{v}}_{t_i + \Delta t}$, the first-order derivative of \mathbf{v}_θ at the timestep t_i can be estimated as: $\mathbf{v}_\theta^{(1)}(\mathbf{Z}_{t_i}, t_i) = \frac{\hat{\mathbf{v}}_{t_i + \Delta t} - \hat{\mathbf{v}}_{t_i}}{\Delta t}$. Substituting this formulation into Eq. (8) results in the practical implementation of the RF-Solver algorithm. The complete sampling process for RF-Solver is presented in Algorithm 1.

Algorithm 1 Sampling process of RF-Solver

Input:

\mathbf{v}_θ ▷ Velocity function
 $t = [t_N, \dots, t_0]$ ▷ Time steps
 $\mathbf{Z}_{t_N} \sim \mathcal{N}(0, I)$ ▷ Initial Gaussian Noise

For $i = N$ **to** 1 **do**

$\Delta t_i \leftarrow \frac{1}{2}(t_{i-1} - t_i)$
 $\hat{\mathbf{v}}_{t_i} \leftarrow \mathbf{v}_\theta(\mathbf{Z}_{t_i}, t_i)$
 $\mathbf{Z}_{t_i + \Delta t_i} \leftarrow \mathbf{Z}_{t_i} + \Delta t_i \hat{\mathbf{v}}_{t_i}$
 $\hat{\mathbf{v}}_{t_i + \Delta t_i} \leftarrow \mathbf{v}_\theta(\mathbf{Z}_{t_i + \Delta t_i}, t_i + \Delta t_i)$
 $\mathbf{v}_{t_i}^{(1)} \leftarrow (\hat{\mathbf{v}}_{t_i + \Delta t_i} - \hat{\mathbf{v}}_{t_i}) / \Delta t_i$ ▷ Calculating the Derivatives
 $\mathbf{Z}_{t_{i-1}} \leftarrow \mathbf{Z}_{t_i} + (t_{i-1} - t_i) \hat{\mathbf{v}}_{t_i} + \frac{1}{2} (t_{i-1} - t_i)^2 \mathbf{v}_{t_i}^{(1)}$

Output: \mathbf{Z}_0

Obtaining the sampling form of RF-Solver, we further derive its inversion form. Inversion maps data back into noise, which reverses the sampling process. Following previous methods for DDIM inversion [13, 59], the ODE process can be directly reversed in the limit of small steps. Based on this assumption, the inversion process of RF-Solver (Eq. (8)) can be directly derived as:

$$\begin{aligned} \tilde{\mathbf{Z}}_{t_{i+1}} &= \tilde{\mathbf{Z}}_{t_i} + (t_{i+1} - t_i) \mathbf{v}_\theta(\tilde{\mathbf{Z}}_{t_i}, t_i) \\ &\quad + \frac{1}{2} (t_{i+1} - t_i)^2 \mathbf{v}_\theta^{(1)}(\tilde{\mathbf{Z}}_{t_i}, t_i), \end{aligned} \quad (9)$$

where $\tilde{\mathbf{Z}}_{t_i}$ and $\tilde{\mathbf{Z}}_{t_{i+1}}$ denotes the latents during inversion. Through the high order expansion, the error of the ODE solution in each timestep is reduced from $\mathcal{O}((h_i)^2)$ to $\mathcal{O}((h_i)^3)$, leading to improved performance, particularly in inversion and reconstruction (see Fig. 2). Beyond inversion and reconstruction, RF-Solver can also be applied to any RF-based model (such as FLUX [2] and OpenSora [3]) for other tasks such as sampling and editing, enhancing performance without requiring additional training.

3.3. RF-Edit

Incorporating higher-order terms enables RF-Solver to significantly reduce errors in the ODE-solving process, thereby

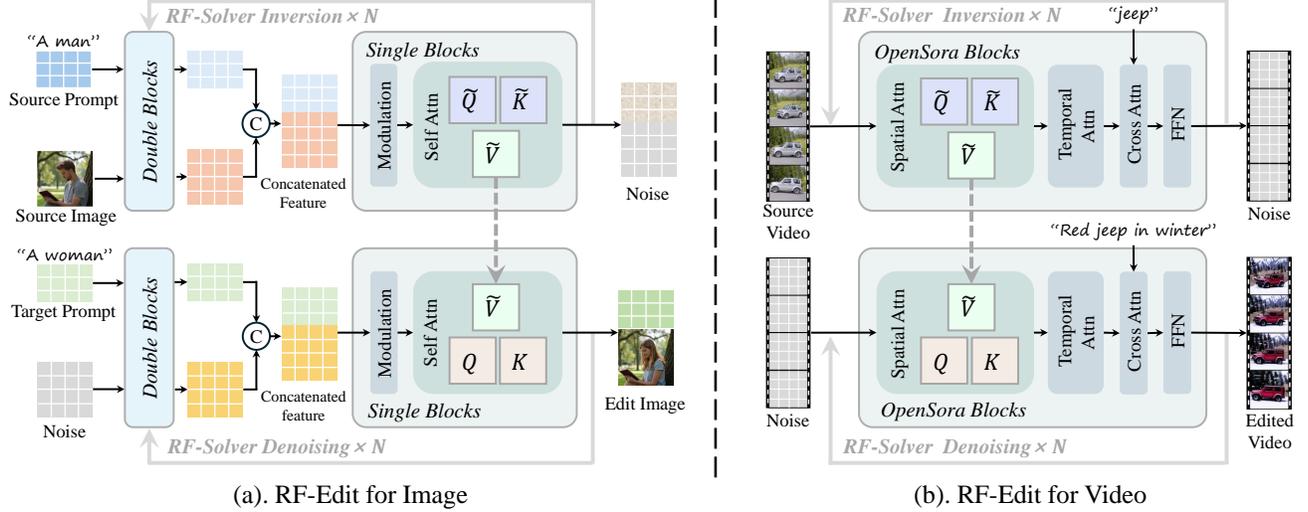


Figure 3. **RF-Edit pipelines for image editing and video editing.** We design two sub-modules for applying RF-Edit to (a). Image editing with FLUX [2] and (b). Video editing with OpenSora [3]. Note that for FLUX, there are multiple Double Blocks, followed by multiple Single Blocks. For OpenSora, there are multiple OpenSora DiT blocks. For simplicity, only one block of each type is depicted in the figure.

enhancing both sampling quality and inversion accuracy. Furthermore, we extend the application of RF-Solver to the more complex real-world image and video editing tasks, which present greater challenges than reconstruction. In such scenarios, preserving the content and structure of the original image is crucial. For example, when replacing an object in a source image with another one, regions unrelated to the object in this image are expected to remain unaffected by the editing process. However, directly applying RF-Solver during the inversion and denoising stages may cause the model to be overly influenced by the target prompt, resulting in unintended modifications in other regions of the source image or video. Similar issues are common across various existing editing methods [19, 57, 63].

To address this problem, we propose RF-Edit, which builds upon the diffusion transformer architecture. Specifically, we focus on the self-attention layer in the last M transformer blocks of v_θ at the last n timesteps during inversion. The self-attention operation can be formulated by:

$$\tilde{\mathbf{F}}_{t_k}^m = \text{Attention}(\tilde{\mathcal{Q}}_{t_k}^m, \tilde{\mathcal{K}}_{t_k}^m, \tilde{\mathcal{V}}_{t_k}^m). \quad (10)$$

Here, $k \in \{N - n, \dots, N\}$, and $m \in \{1, \dots, M\}$, $\tilde{\mathbf{F}}_{t_k}^m$ denotes the output feature of the self-attention module and $\tilde{\mathcal{Q}}_{t_k}^m, \tilde{\mathcal{K}}_{t_k}^m, \tilde{\mathcal{V}}_{t_k}^m$ represent query, key and value for attention during the inversion process, respectively. We extract and store the Value feature $\{\tilde{\mathcal{V}}_{t_k}^m\}$ and $\{\tilde{\mathcal{V}}_{t_k+\Delta t_k}^m\}$ in the process of RF-Solver algorithm (Algorithm 1):

$$\{\tilde{\mathcal{V}}_{t_k}^m\} = \text{Extract}(v_\theta(\tilde{\mathcal{Z}}_{t_k}, t_k)) \quad (11)$$

$$\{\tilde{\mathcal{V}}_{t_k+\Delta t_k}^m\} = \text{Extract}(v_\theta(\tilde{\mathcal{Z}}_{t_k+\Delta t_k}, t_k + \Delta t_k)). \quad (12)$$

During the first n timesteps of denoising, considering the m th transformer block at the timestep k , the original self-

attention can be formulated as:

$$\mathbf{F}_{t_k}^m = \text{Attention}(\mathcal{Q}_{t_k}^m, \mathcal{K}_{t_k}^m, \mathcal{V}_{t_k}^m), \quad (13)$$

where $\mathbf{F}_{t_k}^m$ denotes the output feature of the self-attention module and $\mathcal{Q}_{t_k}^m, \mathcal{K}_{t_k}^m, \mathcal{V}_{t_k}^m$ represent query, key and value for attention during the denoising process, respectively.

In RF-Edit, the above self-attention mechanism is modified to cross-attention where $\mathcal{V}_{t_k}^m$ is replaced by $\tilde{\mathcal{V}}_{t_k}^m$,

$$\mathbf{F}_{t_k}^{m'} = \text{Attention}(\mathcal{Q}_{t_k}^m, \mathcal{K}_{t_k}^m, \tilde{\mathcal{V}}_{t_k}^m). \quad (14)$$

The modified output feature $\mathbf{F}_{t_k}^{m'}$ is then passed to the subsequent modules for further processing.

Similarly, this feature-sharing process is also adopted in the derivative calculation process of RF-Solver:

$$\mathbf{F}_{t_k+\Delta t_k}^{m'} = \text{Attention}(\mathcal{Q}_{t_k+\Delta t_k}^m, \mathcal{K}_{t_k+\Delta t_k}^m, \tilde{\mathcal{V}}_{t_k+\Delta t_k}^m). \quad (15)$$

The proposed RF-Edit framework enables high-quality editing while effectively preserving the structural information of the source image/video. Building on this concept, we design two sub-modules for RF-Edit, specifically tailored for image editing and video editing (Fig. 3). For image editing, RF-Edit employs FLUX [2] as the backbone, which comprises several double blocks and single blocks. Double blocks independently modulate text and image features, while single blocks concatenate these features for unified modulation. In this architecture, RF-Edit shares features within the single blocks, as they capture information from both the source image and the source prompt, enhancing the ability of the model to preserve the structural information of the source image. For video editing, RF-Edit employs OpenSora [3] as the backbone. The DiT blocks in OpenSora include spatial attention, temporal attention, and text cross-attention. Within this architecture, the structural

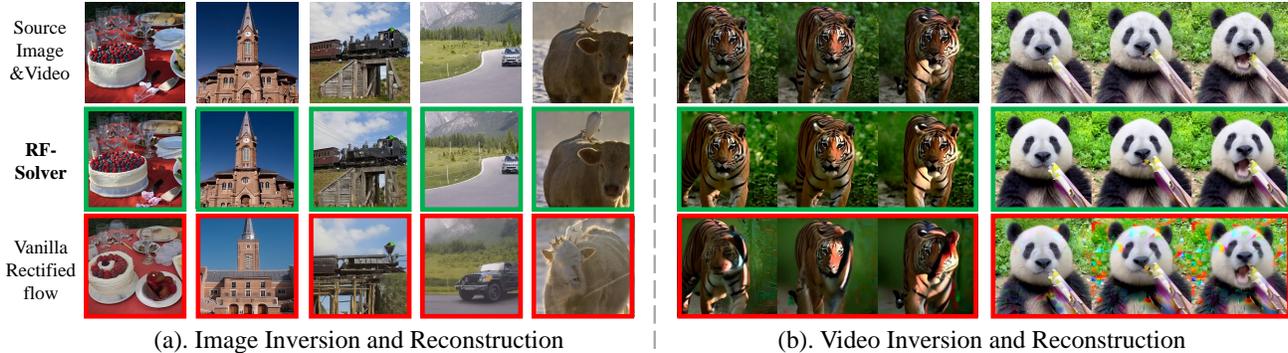


Figure 4. **Qualitative results of image and video reconstruction.** Our method (the second row) demonstrates superior performance compared to the vanilla rectified flow baselines (the third row).

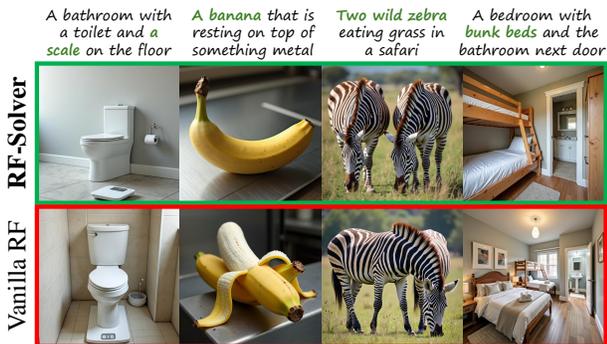


Figure 5. **Qualitative results of text-to-image generation.** By employing the RF-Solver, the model is able to generate images with higher quality (the first row) than baselines (the second row). information of the source video is captured in the spatial attention module, where we implement feature sharing.

4. Experiment

4.1. Setup

Baselines. We select the vanilla Rectified Flow sampler (Euler Sampler) as the primary baseline for all tasks. Besides, for sampling and inversion, we compare our method with DPM-Solver++ [43] and the Heun sampler (which is a 2-order ODE solver). For image editing, we compare our method with P2P [19], DiffEdit [12], SDEdit [44], PnP [63], Pix2pix [49] and RF-Inversion [57]. For video editing tasks, we compare our method with FateZero [52], FLATTEN [10], COVE [67], RAVE [29], Tokenflow [17]. Detailed experimental settings of baselines are provided in the Appendix.

Implementation Details. In the experiment, we adopt the guidance-distilled variant of FLUX [2] for image tasks and OpenSora [3] for video tasks. The derivative computation in RF-Solver requires an additional forward pass, resulting in the network needing to forward twice at each timestep. As a result, when comparing our method with the Rectified Flow baselines, we set the number of timesteps for the vanilla Rectified Flow to be *twice* that of our method to ensure a fair comparison under the same number of function evaluations

	DPMSolver++	RF	RF-Heun	Ours
FID (↓)	24.63	25.33	24.40	24.03
CLIP Score (↑)	30.62	31.01	31.03	31.09

Table 1. **Quantitative results on text-to-image generation.** RF-Solver outperforms several baselines.

(NFE). More information is provided in the Appendix.

Evaluation Metrics For text-to-image sampling, we randomly select 10k images from the MSCOCO validation set [38] and report the FID [20] and CLIP Score [53]. For the inversion and reconstruction task, we report the Mean Square Error (MSE), LPIPS [79], SSIM [71], and PSNR [27]. For image editing tasks, we report the CLIP Score [53] and LPIPS [79]. For video editing tasks, we adopt the metric proposed by [25], including Subject Consistency (SC), Motion Smoothness (MS), Aesthetic Quality (AQ), and Imaging Quality (IQ). Detailed explanations of these metrics are provided in the Appendix.

4.2. Text-to-image Sampling

We compare the performance of our method with DPM-Solver++ [43], the vanilla RF sampler, and Heun sampler on the text-to-image generation task. Both the quantitative (Tab. 1) and qualitative results (Fig. 5) demonstrate the superior performance of RF-Solver in fundamental T2I generation tasks, producing higher-quality images that align more closely with human cognition.

4.3. Inversion and Reconstruction

We conduct experiments on inversion and reconstruction for both image and video modalities, comparing our method with the vanilla RF sampler and the Heun sampler.

Quantitative Comparison. The quantitative comparisons (Tab. 2) are conducted to reflect the similarity between the source and reconstruction results. Our method demonstrates superior performance across all four metrics compared with the vanilla RF sampler and Heun sampler.

Qualitative Comparison. RF-Solver effectively reduces the error in the solution of RF ODE, thereby increasing the accuracy of the reconstruction. As illustrated in Fig. 4(a),

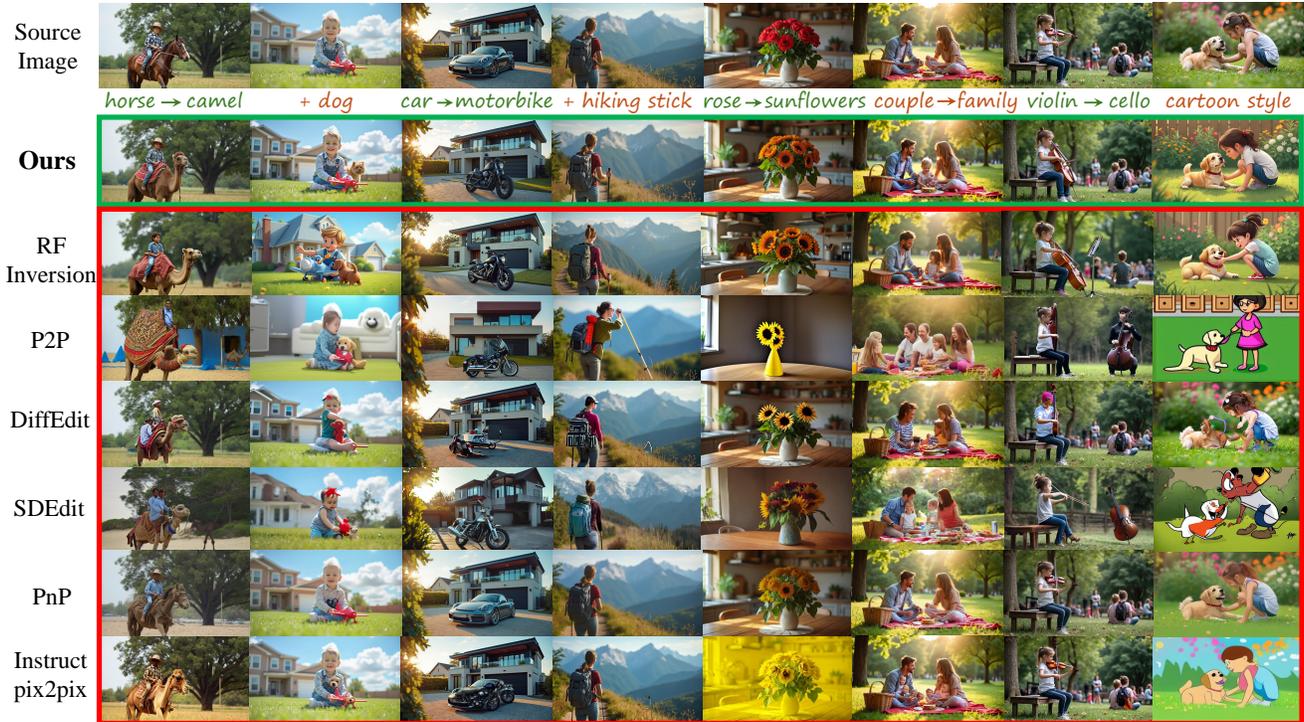


Figure 6. **Qualitative comparison of image editing.** With RF-Solver and feature-sharing mechanism in RF-Edit, our method can successfully handle various kinds of image editing cases, outperforming the previous SOTA methods. Zoom in for the best views.

	Method	MSE (\downarrow)	LPIPS (\downarrow)	SSIM (\uparrow)	PSNR (\uparrow)
image	RF	0.0268	0.6253	0.7626	28.28
	RF-Heun	0.0117	0.4696	0.8924	29.67
	Ours	0.0094	0.4242	0.9271	29.83
video	RF	0.0206	0.4159	0.8134	18.12
	RF-Heun	0.0156	0.3554	0.8711	18.29
	Ours	0.0139	0.3299	0.8805	18.32

Table 2. **Quantitative results on inversion and reconstruction.** Our method significantly improves the accuracy of reconstruction for both images and videos.

the image reconstruction results using vanilla rectified flow exhibit noticeable drift from the source image, with significant alterations to the appearance of subjects in the image. For video reconstruction, as shown in Fig. 4(b), the baseline reconstruction results suffer from distortion. In contrast, RF-Solver significantly alleviates these issues, achieving more satisfactory results.

4.4. Editing

We conduct experiments to evaluate the image and video editing performance of our method. Image editing usually involves replacing the subject in the image with another one, adding new items, and global editing. For the first two types of editing, the background of the source image is expected to remain unchanged after editing. For global editing such as style transfer, the overall structure of the source image is expected to remain unchanged. Recent mainstream video editing methods usually focus on replacing the sub-

	P2P	DiffEdit	SDEdit	PnP	Pix2Pix	RF-Inv	Ours
LPIPS (\downarrow)	0.419	0.157	0.394	0.080	0.155	0.318	0.149
CLIP Score (\uparrow)	30.70	32.68	31.61	30.58	32.33	33.02	33.66

Table 3. **Quantitative results of image editing.** RF-Edit effectively edit the images according to the prompts while preserving the integrity of unrelated regions.

jects and performing global editing for the source video.

Quantitative Comparison. In image editing, Our method outperforms all other methods in CLIP score (Tab. 3), indicating that the edited images align well with the user-provided prompts. For LPIPS, it is noted that PnP [63] has a much lower value than all other methods. Based on the qualitative results (Fig. 6), it can be seen that PnP is only suitable for editing cases that do not significantly modify the structure or shape of the source image (such as changing red roses into yellow sunflowers). It fails in the case of shape editing, resulting in an image very similar to the source. Consequently, although PnP has the lowest LPIPS score, its CLIP score is the lowest.

For video editing, RF-Edit achieves higher scores on the popular VBench [25] metrics (Tab. 4). The results illustrate that our method successfully maintains temporal consistency while demonstrating superior visual quality.

Qualitative Comparison. For image editing, we compare the performance of our method with several baselines across different types of editing tasks (Fig. 6). The baseline methods often suffer from background changes or fail to perform

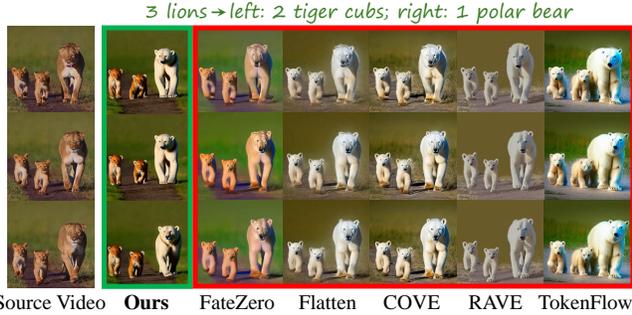


Figure 7. **Qualitative comparison of video editing.** The first video comprises 200 frames with a resolution of 512×512 , while the second video contains 60 frames with a resolution of 1024×768 (frames are compressed for a neat layout in the figure).

	FateZero	Flatten	COVE	RAVE	Tokenflow	Ours
SC (\uparrow)	0.9382	0.9420	0.9433	0.9292	0.9439	0.9501
MS (\uparrow)	0.9611	0.9528	0.9697	0.9519	0.9632	0.9712
AQ (\uparrow)	0.6092	0.6329	0.6717	0.6586	0.6742	0.6796
IQ (\uparrow)	0.6898	0.7024	0.7163	0.6917	0.7128	0.7207

Table 4. **Quantitative results of video editing.** RF-Edit outperforms several previous SOTA video editing methods.

	Metric	RF	RF-Solver-2	RF-Solver-3
Sampling	FID (\downarrow)	25.33	24.03	23.96
	CLIP Score (\uparrow)	31.01	31.09	31.09
Inversion	MSE (\downarrow)	0.0268	0.0094	0.0131
	LPIPS (\downarrow)	0.6253	0.4242	0.4817
Editing	LPIPS (\downarrow)	0.1524	0.1494	0.1503
	CLIP Score (\uparrow)	32.97	33.66	33.18

Table 5. **Ablation study on the Taylor Expansion order.** We select the 2-order expansion (*i.e.* RF-Solver-2) for various downstream tasks due to its effectiveness and simplicity.

the desired edits. In contrast, our method demonstrates satisfying performance, effectively achieving a balanced trade-off between the fidelity to the target prompt and the preservation of the source image.

For video editing, we primarily evaluate the performance of our method on long videos (200 frames) and high-resolution videos (1280×768). The qualitative results are shown in Fig. 7. RF-Edit illustrates impressive performance in *handling complicated editing cases* (*e.g.*, modifying the leftmost lion among three lions into a white polar bear and changing the other two small lions into orange tiger cubs), whereas all other baseline methods fail in this scenario. RF-Edit also demonstrates strong performance in global editing tasks, such as transforming scenes into autumn.

4.5. Ablation Study

We conduct ablation studies to illustrate the effectiveness of RF-Solver and RF-Edit. Without loss of generality, these ablation studies are performed on the image tasks using FLUX [2] as the base model.

Taylor Expansion Order of RF-Solver. We investigated the impact of the Taylor expansion order in RF-Solver (Tab. 5) under the same NFE across different or-

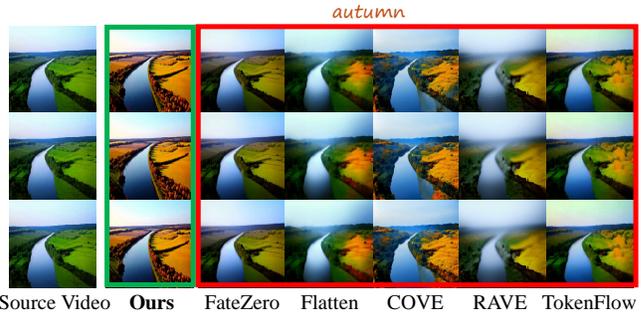


Figure 8. **Ablation study of feature-sharing step in RF-Edit.** A too-small feature-sharing step results in the inconsistency between source and target images. Conversely, a too-large feature-sharing step can lead to the failure of editing.

ders. The second-order expansion demonstrated a significant improvement across various tasks compared to the first-order expansion (*i.e.*, the vanilla rectified flow). However, higher-order expansions do not yield further enhancements. We speculate that this is primarily due to higher-order Taylor expansions requiring more inference steps per timestep. With a fixed NFE, this results in a reduced overall number of timesteps compared to lower-order expansions, leading to suboptimal performance. Moreover, computing the higher-order derivatives of $v_{\theta}(\mathbf{Z}_{t_i}, t_i)$ substantially increases the complexity of the algorithm, posing challenges for practical applications. Consequently, we predominantly employed second-order expansion in our experiments.

Feature Sharing Steps of RF-Edit. RF-Edit leverages feature sharing to maintain the structural consistency between original images and edited images. However, an excessive number of feature-sharing steps may result in the edited output being overly similar to the source image, ultimately undermining the intended editing objectives (Fig. 8). To investigate the impact of feature-sharing steps on editing results, we incrementally increase the number of feature-sharing steps applied to the same image. Due to the varying levels of difficulty that different images presented to the model, the optimal number of sharing steps may differ across cases. Experimental results reveal that setting the sharing step to 5

effectively meets the editing requirements for most images. Additionally, we can customize the sharing step for each image to identify the most satisfying outcome.

5. Conclusion

In this paper, we propose RF-Solver, a versatile sampler for the rectified flow model that solves the rectified flow ODE with reduced error, thus enhancing the image and video generation quality across various tasks such as sampling and reconstruction. Based on RF-Solver, we further propose RF-Edit, which achieves high-quality editing performance while effectively preserving the structural information in source images or videos. Extensive experiments demonstrate the versatility and effectiveness of our method.

References

- [1] Comfyui-fluxtapoz. <https://github.com/logtd/ComfyUI-Fluxtapoz>. 1
- [2] Flux. <https://github.com/black-forest-labs/flux/>. 1, 2, 3, 4, 5, 6, 8
- [3] OpenSora. <https://github.com/hpcaitech/Open-Sora/>. 1, 2, 3, 4, 5, 6
- [4] Omri Avrahami, Ohad Fried, and Dani Lischinski. Blended latent diffusion. *ACM transactions on graphics (TOG)*, 42(4):1–11, 2023. 3
- [5] Manuel Brack, Felix Friedrich, Dominik Hintersdorf, Lukas Struppek, Patrick Schramowski, and Kristian Kersting. Sega: Instructing text-to-image models using semantic guidance. *Advances in Neural Information Processing Systems*, 36:25365–25389, 2023. 3
- [6] Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiaohu Qie, and Yinqiang Zheng. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In *International Conference on Computer Vision*, pages 22560–22570, 2023. 3
- [7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 1
- [8] Wenhao Chai, Xun Guo, Gaoang Wang, and Yan Lu. Stable-video: Text-driven consistency-aware diffusion video editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 23040–23050, 2023. 3
- [9] Weifeng Chen, Yatai Ji, Jie Wu, Hefeng Wu, Pan Xie, Jiashi Li, Xin Xia, Xuefeng Xiao, and Liang Lin. Control-a-video: Controllable text-to-video generation with diffusion models. *arXiv preprint arXiv:2305.13840*, 2023. 3
- [10] Yuren Cong, Mengmeng Xu, Christian Simon, Shoufa Chen, Jiawei Ren, Yanping Xie, Juan-Manuel Perez-Rua, Bodo Rosenhahn, Tao Xiang, and Sen He. Flatten: optical flow-guided attention for consistent text-to-video editing. *arXiv preprint arXiv:2310.05922*, 2023. 3, 6
- [11] Guillaume Couairon, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. Diffedit: Diffusion-based semantic image editing with mask guidance. *arXiv preprint arXiv:2210.11427*, 2022. 3
- [12] Guillaume Couairon, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. Diffedit: Diffusion-based semantic image editing with mask guidance. In *ICLR 2023 (Eleventh International Conference on Learning Representations)*, 2023. 6
- [13] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. 4
- [14] Xiaoyue Duan, Shuhao Cui, Guoliang Kang, Baochang Zhang, Zhengcong Fei, Mingyuan Fan, and Junshi Huang. Tuning-free inversion-enhanced control for consistent image editing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1644–1652, 2024. 2
- [15] Adham Elarabawy, Harish Kamath, and Samuel Denton. Direct inversion: Optimization-free text-driven real image editing with diffusion models. *arXiv preprint arXiv:2211.07825*, 2022. 3
- [16] Xiang Fan, Anand Bhattad, and Ranjay Krishna. Videoshop: Localized semantic video editing with noise-extrapolated diffusion inversion. *arXiv preprint arXiv:2403.14617*, 2024. 2
- [17] Michal Geyer, Omer Bar-Tal, Shai Bagon, and Tali Dekel. Tokenflow: Consistent diffusion features for consistent video editing. *arXiv preprint arXiv:2307.10373*, 2023. 3, 6
- [18] Jiayi Guo, Xingqian Xu, Yifan Pu, Zanlin Ni, Chaofei Wang, Manushree Vasu, Shiji Song, Gao Huang, and Humphrey Shi. Smooth diffusion: Crafting smooth latent spaces in diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7548–7558, 2024. 2
- [19] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 2, 3, 5, 6
- [20] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 6
- [21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2, 3
- [22] Seongmin Hong, Kyeonghyun Lee, Suh Yoon Jeon, Hyewon Bae, and Se Young Chun. On exact inversion of dpm-solvers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7069–7078, 2024. 2
- [23] Nisha Huang, Fan Tang, Weiming Dong, Tong-Yee Lee, and Changsheng Xu. Region-aware diffusion for zero-shot text-driven image editing. *arXiv preprint arXiv:2302.11797*, 2023. 3
- [24] Yi Huang, Jiancheng Huang, Yifan Liu, Mingfu Yan, Jiayi Lv, Jianzhuang Liu, Wei Xiong, He Zhang, Shifeng Chen, and Liangliang Cao. Diffusion model-based image editing: A survey. *arXiv preprint arXiv:2402.17525*, 2024. 3
- [25] Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing Wu, Qingyang Jin,

- Nattapol Chanpaisit, et al. Vbench: Comprehensive benchmark suite for video generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21807–21818, 2024. 6, 7, 1
- [26] Inbar Huberman-Spiegelglas, Vladimir Kulikov, and Tomer Michaeli. An edit friendly ddpm noise space: Inversion and manipulations. In *Conference on Computer Vision and Pattern Recognition*, pages 12469–12478, 2024. 3
- [27] Quan Huynh-Thu and Mohammed Ghanbari. Scope of validity of psnr in image/video quality assessment. *Electronics letters*, 44(13):800–801, 2008. 6
- [28] Xuan Ju, Ailing Zeng, Yuxuan Bian, Shaoteng Liu, and Qiang Xu. Pnp inversion: Boosting diffusion-based editing with 3 lines of code. In *The Twelfth International Conference on Learning Representations*, 2024. 2
- [29] Ozgur Kara, Bariscan Kurtkaya, Hidir Yesiltepe, James M Rehg, and Pinar Yanardag. Rave: Randomized noise shuffling for fast and consistent video editing with diffusion models. In *Conference on Computer Vision and Pattern Recognition*, pages 6507–6516, 2024. 3, 6
- [30] Yoni Kasten, Dolev Ofri, Oliver Wang, and Tali Dekel. Layered neural atlases for consistent video editing. *ACM Transactions on Graphics (TOG)*, 40(6):1–12, 2021. 3
- [31] Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5148–5157, 2021. 2
- [32] Max Ku, Cong Wei, Weiming Ren, Huan Yang, and Wenhu Chen. Anyv2v: A plug-and-play framework for any video-to-video editing tasks. *arXiv preprint arXiv:2403.14468*, 2024. 2
- [33] LAION-AI. aesthetic-predictor. <https://github.com/LAION-AI/aesthetic-predictor>, 2022. 2
- [34] Sangyun Lee, Zinan Lin, and Giulia Fanti. Improving the training of rectified flows. *arXiv preprint arXiv:2405.20320*, 2024. 2
- [35] Yao-Chih Lee, Ji-Ze Genevieve Jang, Yi-Ting Chen, Elizabeth Qiu, and Jia-Bin Huang. Shape-aware text-driven layered video editing. In *Conference on Computer Vision and Pattern Recognition*, pages 14317–14326, 2023. 3
- [36] Shanglin Li, Bohan Zeng, Yutang Feng, Sicheng Gao, Xiuhui Liu, Jiaming Liu, Lin Li, Xu Tang, Yao Hu, Jianzhuang Liu, et al. Zone: Zero-shot instruction-guided local editing. In *Conference on Computer Vision and Pattern Recognition*, pages 6254–6263, 2024. 3
- [37] Zhen Li, Zuo-Liang Zhu, Ling-Hao Han, Qibin Hou, Chun-Le Guo, and Ming-Ming Cheng. Amt: All-pairs multi-field transforms for efficient frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9801–9810, 2023. 2
- [38] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 6
- [39] Qiang Liu. Rectified flow: A marginal preserving approach to optimal transport. *arXiv preprint arXiv:2209.14577*, 2022. 2, 4
- [40] Shaoteng Liu, Yuechen Zhang, Wenbo Li, Zhe Lin, and Jiaya Jia. Video-p2p: Video editing with cross-attention control. In *Conference on Computer Vision and Pattern Recognition*, pages 8599–8608, 2024. 2, 3
- [41] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022. 2
- [42] Xingchao Liu, Chengyue Gong, et al. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *International Conference on Learning Representations*, 2022. 3
- [43] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *arXiv preprint arXiv:2211.01095*, 2022. 3, 6
- [44] 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 *International Conference on Learning Representations*, 2022. 2, 6
- [45] Daiki Miyake, Akihiro Iohara, Yu Saito, and Toshiyuki Tanaka. Negative-prompt inversion: Fast image inversion for editing with text-guided diffusion models. *arXiv preprint arXiv:2305.16807*, 2023. 3
- [46] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images using guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6038–6047, 2023. 2, 3
- [47] Thao Nguyen, Yuheng Li, Utkarsh Ojha, and Yong Jae Lee. Visual instruction inversion: Image editing via image prompting. *Advances in Neural Information Processing Systems*, 36, 2024. 2
- [48] Hao Ouyang, Qiuyu Wang, Yuxi Xiao, Qingyan Bai, Juntao Zhang, Kecheng Zheng, Xiaowei Zhou, Qifeng Chen, and Yujun Shen. Codef: Content deformation fields for temporally consistent video processing. In *Conference on Computer Vision and Pattern Recognition*, pages 8089–8099, 2024. 3
- [49] Gaurav Parmar, Krishna Kumar Singh, Richard Zhang, Yijun Li, Jingwan Lu, and Jun-Yan Zhu. Zero-shot image-to-image translation. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023. 3, 6
- [50] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *International Conference on Computer Vision*, pages 4195–4205, 2023. 2
- [51] Xinyu Peng, Ziyang Zheng, Wenrui Dai, Nuoqian Xiao, Chenglin Li, Junni Zou, and Hongkai Xiong. Improving diffusion models for inverse problems using optimal posterior covariance. In *Forty-first International Conference on Machine Learning*, 2024. 2
- [52] Chenyang Qi, Xiaodong Cun, Yong Zhang, Chenyang Lei, Xintao Wang, Ying Shan, and Qifeng Chen. Fatezero: Fusing attentions for zero-shot text-based video editing. In

- Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15932–15942, 2023. 3, 6
- [53] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 6
- [54] Hareesh Ravi, Sachin Kelkar, Midhun Harikumar, and Ajinkya Kale. Predictor: Text guided image editing with diffusion prior. *arXiv preprint arXiv:2302.07979*, 2023. 3
- [55] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2
- [56] Litu Rout, Yujia Chen, Abhishek Kumar, Constantine Caramanis, Sanjay Shakkottai, and Wen-Sheng Chu. Beyond first-order tweedie: Solving inverse problems using latent diffusion. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 3
- [57] L Rout, Y Chen, N Ruiz, C Caramanis, S Shakkottai, and W Chu. Semantic image inversion and editing using rectified stochastic differential equations. 2024. 3, 5, 6
- [58] Chaehun Shin, Heeseung Kim, Che Hyun Lee, Sang-gil Lee, and Sungroh Yoon. Edit-a-video: Single video editing with object-aware consistency. In *Asian Conference on Machine Learning*, pages 1215–1230. PMLR, 2024. 2
- [59] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. 2, 3, 4
- [60] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. 3
- [61] Wenhao Sun, Rong-Cheng Tu, Jingyi Liao, and Dacheng Tao. Diffusion model-based video editing: A survey. *arXiv preprint arXiv:2407.07111*, 2024. 3
- [62] Haotian Tang, Yecheng Wu, Shang Yang, Enze Xie, Junsong Chen, Junyu Chen, Zhuoyang Zhang, Han Cai, Yao Lu, and Song Han. Hart: Efficient visual generation with hybrid autoregressive transformer. *arXiv preprint arXiv:2410.10812*, 2024. 2
- [63] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. In *Conference on Computer Vision and Pattern Recognition*, pages 1921–1930, 2023. 2, 3, 5, 6, 7
- [64] Bram Wallace, Akash Gokul, and Nikhil Naik. Edict: Exact diffusion inversion via coupled transformations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22532–22541, 2023. 3
- [65] Fangyikang Wang, Hubery Yin, Yuejiang Dong, Huminhao Zhu, Chao Zhang, Hanbin Zhao, Hui Qian, and Chen Li. Belm: Bidirectional explicit linear multi-step sampler for exact inversion in diffusion models. *arXiv preprint arXiv:2410.07273*, 2024. 2, 3
- [66] Fu-Yun Wang, Ling Yang, Zhaoyang Huang, Mengdi Wang, and Hongsheng Li. Rectified diffusion: Straightness is not your need in rectified flow. *arXiv preprint arXiv:2410.07303*, 2024. 2, 4
- [67] Jiangshan Wang, Yue Ma, Jiayi Guo, Yicheng Xiao, Gao Huang, and Xiu Li. Cove: Unleashing the diffusion feature correspondence for consistent video editing. *arXiv preprint arXiv:2406.08850*, 2024. 3, 6
- [68] Qian Wang, Biao Zhang, Michael Birsak, and Peter Wonka. Instructedit: Improving automatic masks for diffusion-based image editing with user instructions. *arXiv preprint arXiv:2305.18047*, 2023. 3
- [69] Wen Wang, Yan Jiang, Kangyang Xie, Zide Liu, Hao Chen, Yue Cao, Xinlong Wang, and Chunhua Shen. Zero-shot video editing using off-the-shelf image diffusion models. *arXiv preprint arXiv:2303.17599*, 2023. 3
- [70] Yuhan Wang, Liming Jiang, and Chen Change Loy. Styleinv: A temporal style modulated inversion network for unconditional video generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22851–22861, 2023. 2
- [71] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6
- [72] Enze Xie, Junsong Chen, Junyu Chen, Han Cai, Yujun Lin, Zhekai Zhang, Muyang Li, Yao Lu, and Song Han. Sana: Efficient high-resolution image synthesis with linear diffusion transformers. *arXiv preprint arXiv:2410.10629*, 2024. 2
- [73] Shuai Yang, Yifan Zhou, Ziwei Liu, and Chen Change Loy. Rerender a video: Zero-shot text-guided video-to-video translation. In *SIGGRAPH Asia 2023 Conference Papers*, pages 1–11, 2023. 3
- [74] Shuai Yang, Yifan Zhou, Ziwei Liu, and Chen Change Loy. Fresco: Spatial-temporal correspondence for zero-shot video translation. In *Conference on Computer Vision and Pattern Recognition*, pages 8703–8712, 2024. 3
- [75] Xiaofeng Yang, Cheng Chen, Xulei Yang, Fayao Liu, and Guosheng Lin. Text-to-image rectified flow as plug-and-play priors. *arXiv preprint arXiv:2406.03293*, 2024. 3
- [76] Zhen Yang, Ganggui Ding, Wen Wang, Hao Chen, Bohan Zhuang, and Chunhua Shen. Object-aware inversion and reassembly for image editing. *arXiv preprint arXiv:2310.12149*, 2023. 3
- [77] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024. 2
- [78] Guoqiang Zhang, Jonathan P Lewis, and W Bastiaan Kleijn. Exact diffusion inversion via bidirectional integration approximation. In *European Conference on Computer Vision*, pages 19–36. Springer, 2024. 2
- [79] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 6

- [80] Yabo Zhang, Yuxiang Wei, Dongsheng Jiang, Xiaopeng Zhang, Wangmeng Zuo, and Qi Tian. Controlvideo: Training-free controllable text-to-video generation. *arXiv preprint arXiv:2305.13077*, 2023. 3

Taming Rectified Flow for Inversion and Editing

Supplementary Material

6. Experimental Settings

6.1. Baselines and Implementation Details

Text-to-Image Generation. We compare our methods with the following baselines: FLUX with the vanilla sampler, Heun Solver, and DPM-Solver. The Heun Solver is a second-order ODE solver that can be applied to pretrained rectified flow to solve the ODE more precisely. DPM-Solver is a high-order sampler for diffusion ODE, which is not suitable for RF-based models like FLUX. As an alternative, we apply the DPM-Solver on Stable Diffusion to evaluate its performance. For FLUX with the vanilla sampler and the Heun Solver, we randomly select 10000 images from the MS-COCO validation dataset and use their caption as the prompt for generation. The resolution of generated images is 1024×1024 . For DPM-Solver, we adopt the implementation from the diffuser, adopting its default setting to generate images. The total NFE for generating one image is set to 10 for both our method and baselines.

Inversion. We compare the performance of our methods among RF with the vanilla sampler and the Heun sampler. For image inversion, we also use the images from the MS-COCO validation set. For video inversion, we select videos from social media platforms such as TikTok and other publicly available sources. We have observed the quality of the text prompts significantly influence the quality of inversion. Consequently, we employ GPT-4o to generate detailed captions for both images and videos, which are then used in the inversion tasks. The total NFE for generating one image/video is set to 50 for both our method and baselines.

Editing. For image editing, we compare our methods with RF-inversion and several diffusion-based editing methods. For RF-inversion, we adopt the implementation in ComfyUI [1]. For other baselines, we use their implementation from diffusers. For each baseline, we adjust the relevant hyper-parameters to achieve optimal results. For video editing, we use the official codes of all the baseline methods and tune the hyper-parameters to achieve satisfactory results. For image editing, we share the features of the last 19 single blocks in FLUX. For video editing, we share the features of the last 14 blocks in Open-Sora. We adjust the hyper-parameter of feature-sharing steps to achieve better results for both image and video editing.

6.2. Evaluation Metrics

For text-to-image sampling, we report Fréchet Inception Distance (FID) and CLIP Scores. The FID is a metric used to evaluate the quality of generated images by assessing the similarity between the distributions of real and generated

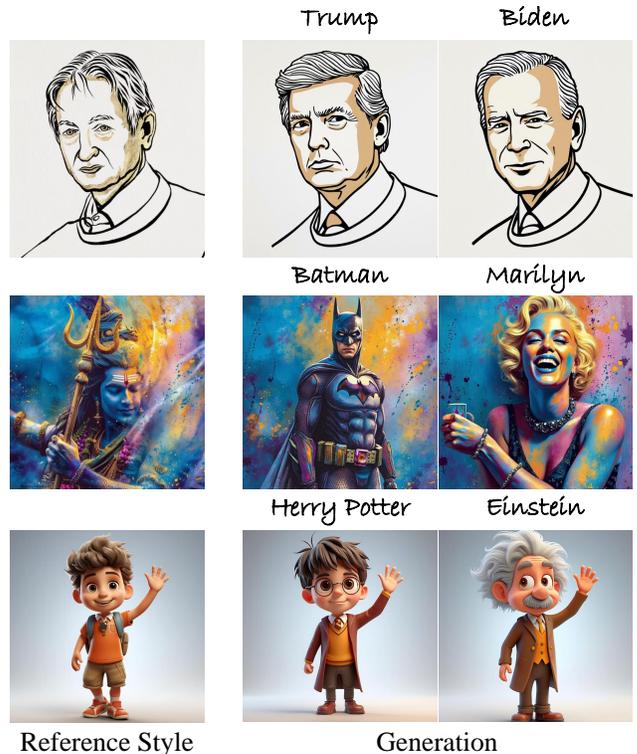


Figure 9. Stylization Results.

image features, typically extracted using a pre-trained Inception network. The CLIP Score evaluates the alignment between generated images and textual descriptions by measuring the similarity of their embeddings within a shared multimodal space using the CLIP model.

For Inversion tasks, our evaluation metrics include MSE, LPIPS, SSIM, and PSNR. MSE measures the average squared difference between predicted and ground-truth values, quantifying the overall error in pixel intensity. LPIPS assesses perceptual similarity between images by comparing deep feature representations extracted from neural networks, aligning with human perception. SSIM evaluates image quality by comparing luminance, contrast, and structure to measure the similarity between the reference and reconstructed images. PSNR quantifies the ratio between the maximum possible signal value and the power of noise, commonly used to assess image reconstruction quality.

For video editing, we adopt the VBench Metrics [25]. The evaluation criteria include Subject Consistency, Motion Smoothness, Aesthetic Quality, and Imaging Quality. Subject Consistency measures whether the subject (*e.g.*, a person) remains consistent throughout the video by computing the similarity of DINO features [7] across frames. Motion

Smoothness assesses the smoothness of motion in the generated video using motion priors from the video frame interpolation model [37]. Aesthetic Quality evaluates the artistic and visual appeal of each frame as perceived by humans, leveraging the LAION aesthetic predictor [33]. Imaging Quality examines the level of distortion in the generated frames (*e.g.*, blurring or flickering) based on the MUSIQ image quality predictor [31].

7. Stylization

We provide more results on image stylization (Fig. 9). In this task, we perform inversion on the image with a reference style, following the same overall pipeline as in image editing. RF-Edit demonstrates superior performance in stylized generation, effectively preserving the style of the source image while generating images that align with the user-provided prompts.

8. Limitation

Despite achieving impressive performance across various tasks, our methods occasionally exhibit performance instability, particularly in video editing. Additionally, the FLUX and OpenSora models contain a large number of parameters, leading to significant memory consumption, which may limit the applicability of our method on resource-constrained devices. To mitigate these issues, we plan to further optimize our method and explore its efficacy on more lightweight models.