

Unveil Inversion and Invariance in Flow Transformer for Versatile Image Editing

Pengcheng Xu^{1,2} Boyuan Jiang² Xiaobin Hu² Donghao Luo² Qingdong He²
 Jiangning Zhang² Chengjie Wang² Yunsheng Wu² Charles Ling¹ Boyu Wang¹

¹Western University ²Tencent

<https://pengchengpcx.github.io/EditFT/>

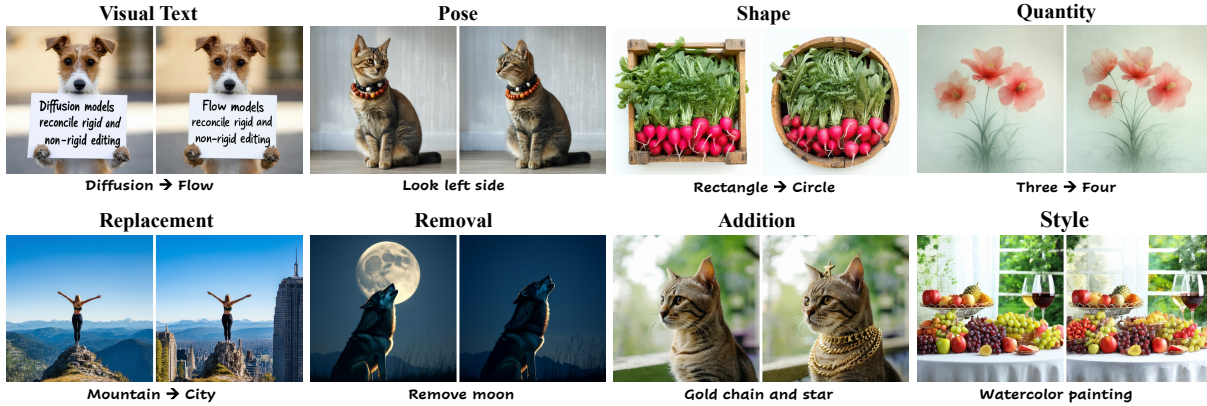


Figure 1. Our framework reconciles the invariance control for rigid and non-rigid editing, enabling versatile editing via flow transformer.

Abstract

Leveraging the large generative prior of the flow transformer for tuning-free image editing requires authentic inversion to project the image into the model’s domain and a flexible invariance control mechanism to preserve non-target contents. However, the prevailing diffusion inversion performs deficiently in flow-based models, and the invariance control cannot reconcile diverse rigid and non-rigid editing tasks. To address these, we systematically analyze the **inversion and invariance** control based on the flow transformer. Specifically, we unveil that the Euler inversion shares a similar structure to DDIM yet is more susceptible to the approximation error. Thus, we propose a two-stage inversion to first refine the velocity estimation and then compensate for the leftover error, which pivots closely to the model prior and benefits editing. Meanwhile, we propose the invariance control that manipulates the text features within the adaptive layer normalization, connecting the changes in the text prompt to image semantics. This mechanism can simultaneously preserve the non-target contents while allowing rigid and non-rigid manipulation, enabling a wide range of editing types such as visual text, quantity, facial expression, etc. Experiments on versatile scenarios validate that our framework achieves flexible and accurate editing, unlocking the potential of the flow transformer for versatile image editing.

1. Introduction

The success of the diffusion model (DM) in large-scale text-to-image generation [33, 35, 37] enables editing images with flexible text-guided interface and rich expressiveness [11, 18, 28, 30]. Image editing requires the model conditioning on the input image. The instruction-based methods retrain the text-to-image (T2I) DM by the triplet data of the original-and-edited images, and the editing instruction [4, 8]. However, manually crafting such triples is not scalable and may cause bias [51]. In contrast, tuning-free methods leverage the model itself to implicitly construct the input image condition, which eliminates additional training and ensures scalability with the large T2I model. This approach has recently garnered significant research interest. Despite this, most studies focus on DM editing, while the potential of recent flow-based T2I models (FMs), which incorporate transformer architecture (i.e., DiT [31]), richer data priors, and better text-to-image alignment [6], remains underexplored.

In this paper, we investigate tuning-free image editing within the context of the flow transformer from two key aspects: **inversion** and **invariance control**. The inversion process projects the image into the model’s domain as an initial noisy latent, while invariance control determines which contents from the original image should be preserved during the regeneration of the edited image. The principle of inversion aims to enhance the faithfulness of the recov-

ered original image and editing ability. Although inversions based on the DDIM [40] in DM have been extensively studied, they face significant challenges with the Euler sampler in flow models (FM). Regarding the invariance control, the attention manipulation techniques [11, 30, 42] in U-Net dominate DM but are deficient at reconciling rigid and non-rigid editing such as altering numbers, movements, and shapes [5, 19, 47]. For instance, preserving non-target contents can hinder changes to the layout and pose of objects. Complementary solutions often require additional diffusion branches and specific refinement [2, 19]. Thus, a reconciled and efficient solution is beneficial for editing.

Based on the above, the editing challenges of the flow transformer stem from two major issues: **inversion for flow** and **invariance control based on the flow transformer**. First, the off-the-shelf diffusion inversions are not optimized for the Euler sampler in flow matching and suffer fallback in practice (Sec:4.1). Developing faithful and editable inversion for flow matching is essential. Second, how to preserve the invariance (i.e. the unedited contents) based on transformer architecture is crucial for tuning-free editing yet underexplored due to the network architecture difference (DiT v.s. U-Net). Specifically, MM-DiT [6] does not have cross-attention as the conditioning mechanism, and simply extending self-attention methods to MM-DiT yields unsatisfied editing. Therefore, effective invariance control must reconcile rigid and non-rigid editing to take full advantage of the rich text-to-image priors for versatile editing.

To address the aforementioned issues, we systematically investigate flow inversion with the Euler method and invariance modeling in MM-DiT. For inversion, we reveal that the reverse ordinary differential equation (ODE) with the Euler method has a similar structure to DDIM inversion but is more affected by the approximation error between two consecutive states. This error makes Euler inversion diverge more from the original image than DDIM inversion in diffusion and significantly deteriorates the editing ability even if the image can be recovered by optimization. So, we proposed a two-stage flow inversion. First, we use the fixed-iteration technique to reduce the approximation error when estimating the flow velocity and get a basic inversion with the Euler method. This makes inversion approximates the authentic generation process. Then based on such inversion, we only need to add mild compensation at each denoising step to recover the exact original image. Since the compensation at each denoising step is small, the inversion trajectory is close to the model’s original domain, thereby preserving the editing ability.

In the investigation of the invariance control in MM-DiT, our key finding is that the replacement of the text features corresponding to the unedited prompt within the adaptive layer normalization (AdaLN) can simultaneously preserve the non-target contents while allowing rigid and non-rigid

editing such as altering layout, quantity, and pose. This contrasts with the self-attention mechanism which tends to be indiscriminative to edited and unedited words in the text prompt and indistinctively injects the contents of the original image into the target, hindering some non-rigid editing types. Consequently, our method can connect changes in the text prompt to the changes in AdaLN features and then guide image generation. Besides, the MM-DiT only has the self-attention that joins the text and image modalities together. This further makes it difficult to disentangle the target edited semantics from the invariance controlled by the attention mechanism and influence the desired editing (see Figure 5). Thus, we choose the AdaLN as the main invariance control in MM-DiT, leaving the attention as a subtle complementary.

In summary, our contributions and findings include:

- We unveil the deficiency and relation of Euler to DDIM inversion and propose a two-stage inversion to reduce the velocity approximation error and recover the exact image while enhancing the editing ability by making the inversion pivoting around the authentic generation process.
- We introduce a flexible invariance control mechanism for the flow transformer (MM-DiT) based on adaptive layer normalization (AdaLN), which reconciles rigid and non-rigid editing and enables versatile editing types.
- Experiments on various editing types validate that our method gives full play to the large prior of the flow transformer for versatile tuning-free image editing.

2. Related Work

Image Editing via Diffusion Models. Leveraging the large prior of the text-to-image diffusion model, earlier editing methods fine-tune the model with the original image and the text prompt [18]. Later approaches mainly fall into three categories. Truncation methods analyze latent noise predictions of the model [2, 19] and filter out undesired editing parts to preserve the non-target regions. The instruction-training methods [4, 8, 10, 26, 38, 51] require collecting a triplet of original and edited images and the editing instruction to re-train the diffusion model to be an editing model, which can be difficult to scale up. The inversion-based methods [5, 11, 30, 42, 48] are tuning-free and first convert the original image into the initial noisy latent and regenerate the edited image with the edited prompt while preserving the invariance (i.e., unchanged content from the original image). Others further improve the applicable ability by combining the large language model [15, 20] and enrich the editing forms [21, 24, 39, 49]. Hu et al. [14] research the latent space for flow-based image editing but the backbone is U-ViT yet the large-scale DiT model, and does not discuss the flow inversion.

Inversion of Diffusion for Real Image Editing. The inversion that converts the image into the initial noisy la-

tent is pivotal in tuning-free editing. One branch leverages DDIM [40] to invert the ODE process. However, the classifier-free guidance (CFG) in editing deteriorates the accumulated error in reversing ODE. Various mitigation solutions explore optimizing the null-text embedding [28], adding corrections [17], better estimation of latents [7], and negative prompt [27]. Another branch adopts the stochastic differential equation (SDE) instead of ODE to mitigate the error by manipulating random noise [3, 16]. Besides these heuristic solutions, others explore the mathematical exact inversion by reformulating the sampling process [13, 43, 44]. Besides these nonparametric solutions, TurboEdit [45] learns the inversion process with a neural network to reduce the inversion time. Unlike these diffusion-oriented inversions, we specifically unveil the deficiency of flow inversion and study efficient inversion based on the velocity field of flow models.

Invariance control in diffusion-based image editing. The invariance control preserves the original image’s unedited contents during editing. The instruction-based methods explicitly add the original image features as the condition and retrain the T2I model into a text-guided image-to-image (TI2I) model [4, 8]. In the tuning-free paradigm, attention-based methods [11, 42] inject attention maps from the original image to preserve the original contents but may struggle with non-rigid editing such as layout and pose. Later methods [5, 47] propose the mutual-self attention that injects K and V values in the later diffusion process to adapt to the layout change but this deteriorates the fidelity for large changes. The refinement approach [2, 19] filters out components of the predicted noise corresponding to the non-target regions to preserve the non-target regions. However, these methods require careful tuning filter for different editing types. Moreover, these approaches are mainly based on diffusion and U-Net. In contrast, to take full advantage of the generation ability of the flow transformer, we systematically investigate inversion for rectified flow and flexible invariance control based on MM-DiT to support rigid and non-rigid editing, enabling versatile editing types.

3. Preliminary

We aim to design a tuning-free framework based on the flow transformer for versatile image editing. In the following, we first briefly introduce the rectified flow and DDIM inversion in diffusion in Sec 3. Then, we analyze the deficiency of Euler inversion in rectified flow and propose the two-stage flow inversion in Sec 4. In Sec 5, we present the image invariance investigation in MM-DiT centered on text features in AdaLN. Last, we summarize the full algorithm.

Rectified Flow. Similar to the diffusion model [12, 41] that learns the probability paths between two distributions, rectified flow [6, 23] linearly interpolates the probability path between two observed distributions $\mathbf{x}_0 \sim p_0$ and $\mathbf{x}_1 \sim p_1$

in Eq. 1 and learns such straight probability transport paths with ODE in Eq. 2 where the transport is modeled by a time-aware velocity $v_\theta(\mathbf{x}_t, t)$.

$$\mathbf{x}_t = t\mathbf{x}_1 + (1-t)\mathbf{x}_0, \quad t \in [0, 1]. \quad (1)$$

$$d\mathbf{x}_t = v_\theta(\mathbf{x}_t, t)dt, \quad \mathbf{x}_0 \sim p_0, \quad t \in [0, 1]. \quad (2)$$

In practice, the velocity v_θ of rectified flow is parameterized by the weights θ of a neural network and can be straightforwardly derived from differentiating Eq. 1 as $\mathbf{x}_1 - \mathbf{x}_0$. The training objective is to directly regress the velocity field in with the least squares loss Eq. 3 and \mathbf{x}_t in Eq. 1.

$$\mathcal{L} = \mathbb{E}_{t \sim \mathcal{U}[0,1], \mathbf{x}_1 \sim p_1} \left[\|\mathbf{x}_1 - \mathbf{x}_0 - v_\theta(\mathbf{x}_t, t)\|^2 \right]. \quad (3)$$

Generally, p_0 is chosen as the standard Gaussian $\mathcal{N}(0, I)$, and p_1 is the target data distribution. The sampling process is to start as \mathbf{x}_0 and integrate the ODE from $t : 0 \rightarrow 1$ to yield a data sample \mathbf{x}_1 , which is practically implemented by solving the discrete ODE by Euler method in Eq. 4 where σ_t is the discrete timestep.

$$\mathbf{x}_{t+1} = \mathbf{x}_t + (\sigma_{t+1} - \sigma_t)v_\theta(\mathbf{x}_t, t) \quad (4)$$

DDIM. DDIM [40] is the implicit probabilistic model trained with DDPM [12] objective and has the corresponding ODE formulation below where \mathbf{x}_t is also the noisy data sample, $\mathbf{s}_\theta(\mathbf{x}_t, t)$ is the noise prediction, $\mathbf{f}_\theta(\mathbf{x}_t, t)$ is the predicted original data sample.

$$\mathbf{x}_{t+1} = \sqrt{\alpha_{t+1}}\mathbf{f}_\theta(\mathbf{x}_t, t) + \sqrt{1 - \alpha_{t+1}}\mathbf{s}_\theta(\mathbf{x}_t, t) \quad (5)$$

$$\mathbf{f}_\theta(\mathbf{x}_t, t) = \frac{\mathbf{x}_t - \sqrt{1 - \alpha_t}\mathbf{s}_\theta(\mathbf{x}_t, t)}{\sqrt{\alpha_t}} \quad (6)$$

4. Inversion for Rectified Flow

The inversion of the sampling process from the data sample \mathbf{x}_1 back to \mathbf{x}_0 can be solved by reverse ODE. In rectified flow, the **exact** Euler inversion is Eq. 7. However, since we move backward, we only have \mathbf{x}_{t+1} but not \mathbf{x}_t . So, we approximate \mathbf{x}_t with \mathbf{x}_{t+1} in Eq. 8. This introduces the **approximation error** at every step and accumulates to deviate the initial latent \mathbf{x}_0 . The same applies to DDIM inversion.

$$\mathbf{x}_t = \mathbf{x}_{t+1} + (\sigma_t - \sigma_{t+1})v_\theta(\mathbf{x}_t, t) \quad (7)$$

$$\approx \mathbf{x}_{t+1} + (\sigma_t - \sigma_{t+1})v_\theta(\mathbf{x}_{t+1}, t) \quad (8)$$

4.1. Deficiency and Relation of Euler to DDIM

Essentially, if we rescale \mathbf{x}_t with $\frac{\mathbf{x}_t}{\sqrt{\alpha_t}}$ and rewrite Eq. 5, DDIM can also be treated as the first-order ODE solver as Euler above, and there has been extensive research on optimizing the DDIM inversion in DM. However, despite the similar formulation of Euler and DDIM inversion, the performance of Euler inversion is significantly worse than the

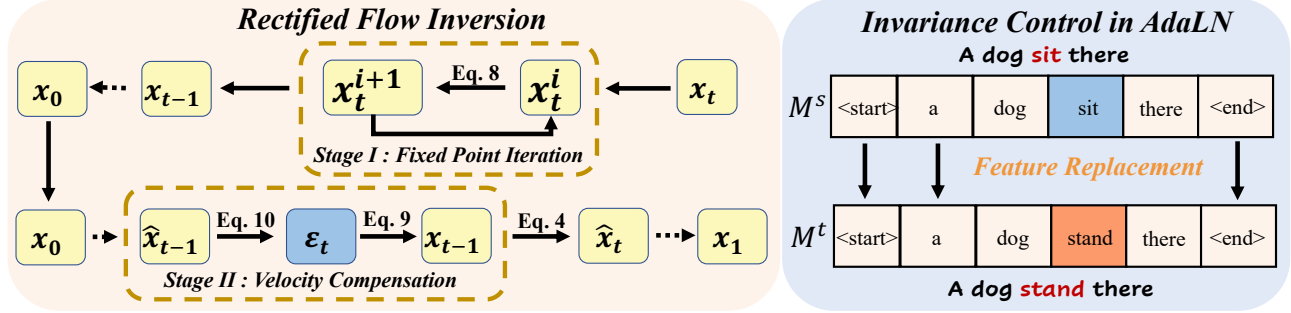


Figure 2. **Framework of rectified flow inversion and editing.** **Left:** The two-stage inversion for rectified flow. A basic inversion trajectory is first constructed to resemble the generation process, and then mild compensation is added to velocity to recover the image exactly. **Right:** Invariance control with text feature replacement within AdaLN during sampling. In the target branch, the unchanged text features are replaced with features from the original image while the edited text features remain intact.

DDIM even if the Euler inversion is evaluated on the more powerful Stable Diffusion 3.5 v.s SD 1.5 as shown in Figure 3. Such phenomena are not expected and also make optimization methods [3, 17, 28] in diffusion invalid for flow inversion.




Figure 3. **Comparison of Euler inversion for rectified flow and DDIM inversion.** With the same inversion steps as 50, the Euler is evaluated on SD3.5 while the DDIM is evaluated on SD1.5.

We argue this is due to the fact that the Euler inversion for rectified flow transformer suffers more from the approximation error. To reduce it, we need to mitigate the gap between $v_\theta(\mathbf{x}_{t+1}, t)$ and $v_\theta(\mathbf{x}_t, t)$. Ideally, we expect to use the exact inversion in Eq. 7 to get the initial noisy latent \mathbf{x}_0 and the transition trajectory $\{\mathbf{x}_1, \dots, \mathbf{x}_t\}$. Such an inversion trajectory is close to the trajectory in the generating process and requires minimal optimization. Thus, it fits the prior distribution of the model and has good editing ability.

4.2. Two-stage Flow Inversion

Our principle is to find an initial latent and inversion trajectory that is as close to the generation process as possible. This has two-fold benefits: 1. **Editing friendly.** The latent and trajectory fit into the prior distribution of the model so that the image can be easily manipulated by changing the text prompt. 2. **Easy to preserve invariance and fidelity.** If the inversion trajectory deviates from the model prior, it requires more effort such as self-attention injection to preserve the invariance and fidelity. However, injecting too much self-attention from the original image can hinder the editing. Thus, we propose a two-stage flow inversion. First, we construct a pivotal inversion close to the generating process. Second, we manually add compensation at

every timestep to eliminate the left mild error.

Stage I: Fixed point iteration with stable velocity. We aim to use the Eq. 7 to get the exact \mathbf{x}_t without the approximation. To achieve this, we need to improve the estimation of $v_\theta(\mathbf{x}_t, t)$ instead of using the approximated $v_\theta(\mathbf{x}_{t+1}, t)$. Note that the input to $v_\theta(\mathbf{x}_t, t)$ and output of Eq. 7 are both \mathbf{x}_t . Inspired by the fixed-point technique [7, 9], we can iteratively apply Eq. 7 to \mathbf{x}_t itself and better estimate the $v_\theta(\mathbf{x}_t, t)$. Concretely, we first initialize \mathbf{x}_t^1 with \mathbf{x}_{t+1} and iteratively apply the following equation to get a series estimation of $\{\mathbf{x}_t^i\}_{i=1}^I$.

$$\mathbf{x}_t^{i+1} = \mathbf{x}_{t+1} + (\sigma_t - \sigma_{t+1})v_\theta(\mathbf{x}_t^i, t) \quad (9)$$

Besides, ideally, the velocity v_θ is $\mathbf{x}_1 - \mathbf{x}_0$ according to Eq. 3 that is a stable constant. This also strengthens the motivation of fixed-point iteration, and we can use this property to average the velocities to get a more stable velocity prediction. In practice, this is also equivalent to averaging the latent series $\{\mathbf{x}_t^i\}_{i=1}^I$.

Stage II: Velocity compensation in editing. In Stage I, we get an inversion trajectory $\{\mathbf{x}_i\}_{i=1}^T$ that significantly reduces the approximation error and is close to the generation process. However, since the fixed-point iteration is a numerical method, the error may still exist when using the inverted \mathbf{x}_0 to recover \mathbf{x}_1 . To exactly recover the original image \mathbf{x}_1 with the inversion trajectory $\{\mathbf{x}_i\}_{i=1}^T$, at every timestep, we manually calculate and add compensation ϵ_t for the velocity during the generating process (editing process) so that the recovery exactly follow the inversion trajectory.

$$\hat{\mathbf{x}}_{t+1} = \mathbf{x}_t + (\sigma_{t+1} - \sigma_t)v_\theta(\mathbf{x}_t, t) \quad (10)$$

$$\epsilon_t = \mathbf{x}_{t+1} - \hat{\mathbf{x}}_{t+1} \quad (11)$$

Note that this stage is done during editing but not traversing the denoising process again. The algorithm is summarized in Algorithm 1. In practice, the inversion pivots around the authentic generation process with mild deviation, which recovers the exact image and provides good editing ability.

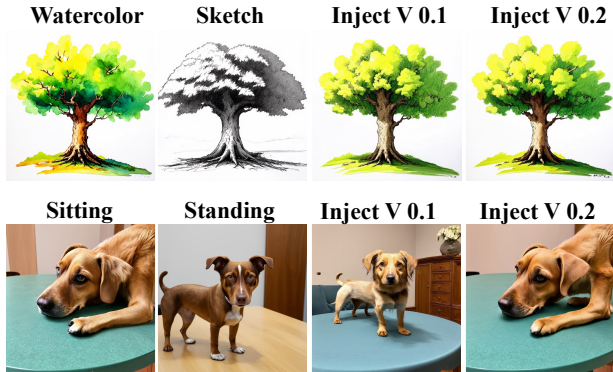


Figure 5. **Invariance control of self-attention with different time periods.** We inject different components (Q,K,V) of self-attention with different time periods to verify the ability to preserve the unedited contents (see Appendix). Here we show injecting V value more than 20% time steps hinders the editing.

6. Experiments

In this section, we systematically validate our methods in various editing types and scenarios. We also ablate the effects of two-stage inversion and iteration steps. We provide both qualitative and quantitative results. Please also refer to the Appendix for more results.

6.1. Setup

Baselines and implementations. We mainly compare our methods with previous state-of-the-art turning-free image editing methods: P2P [11], PnP [42], MasaCtrl [5], InfEdit [47]. All these methods use attention maps in U-Net to control the invariance for editing. Specifically, MasaCtrl and InfEdit address the issue of non-rigid editing. Additionally, we also include the instruction-based editing method InsP2P [4] which is retrained with the triplet of before-and-after image pairs and the editing instruction. For all methods, we use their official implementations for evaluation. We implement our method based on the Stable Diffusion 3.5 and use the basic Euler sampler of rectified flow for sampling and inversion. We use 30 steps for inversion and choose the CFG as 1 and 2 for inversion and editing respectively. We set the fixed-point iteration number as 3. For methods based on diffusion U-Net and DDIM sampling, we use 50 as inversion steps and choose CFG as 7, which follows their default settings.

Evaluation Datasets and Metrics. We evaluate our method based on the PIE benchmark [17] which includes 700 natural and artificial images in total. The dataset covers a wide range of editing types and provides detailed metrics to evaluate both the editing ability and invariance preservation ability at the same time. Concretely, for invariance preservation, the metrics include PSNR, LPIPS, MSE, and SSIM. Note that the numbers of these metrics reported in this paper are scaled. For the editing ability, the metric cal-

culates the CLIP similarity between the whole image and the target prompt (Whole) and between the target regions and the target text prompt (Edited).

6.2. Comparison with previous editing methods

Real image editing and fidelity. We compare our method with others in real image editing. As shown in Figure 6, we evaluate methods with different editing types that represent different levels of geometric changes. Our method can better preserve non-target content and achieve high fidelity. This validates the generalization ability of our method in real image editing. In contrast, methods that use attention injection to preserve the invariance cannot reconcile different editing types. Specifically, P2P and PnP utilize the cross and self-attention injection to preserve better structure information but they also sacrifice the ability to edit the layout and pose of objects. MasaCtrl and InfEdit inject the KV in the later diffusion steps to allow editing of the layout and pose since the structure information tends to form in the early steps. However, their fidelity is affected when facing large changes. For example, MasaCtrl is limited in object replacement. The instruction-based InsP2P is flexible to all editing types but inferior in preserving the structure information since the training data are difficult to construct and may not be accurate, which influences editing fidelity. Please also see the Appendix for more qualitative results.

Table 1. **Quantitative comparisons in real image editing.** Evaluated using the PIE benchmark. Different metrics are scaled.

| Method | Structure | | Background Preservation | | | CLIP Similarity | |
|----------|--------------|--------------|-------------------------|--------------|--------------|-----------------|--------------|
| | Distance ↓ | PSNR ↑ | LPIPS ↓ | MSE ↓ | SSIM ↑ | Whole ↑ | Edited ↑ |
| PnP | 24.29 | 22.46 | 106.06 | 80.45 | 79.68 | <u>25.41</u> | 22.62 |
| P2P | 13.44 | 27.03 | 60.67 | 35.86 | 84.11 | 24.75 | 21.86 |
| MasaCtrl | 24.70 | 22.64 | 87.94 | 81.09 | 81.33 | 24.38 | 21.35 |
| InsP2P | 57.91 | 20.82 | 158.63 | 227.78 | 76.26 | 23.61 | 21.64 |
| InfEdit | 24.70 | 26.31 | 87.94 | 75.19 | 81.33 | 23.67 | 21.86 |
| Ours | <u>18.17</u> | <u>26.62</u> | <u>80.55</u> | <u>40.24</u> | 91.50 | 25.74 | <u>22.27</u> |

Quantitative comparison. The quantitative results are summarized in Table 1. The comparison methods are implemented with the most suitable diffusion-based inversions for high performance. P2P and MasaCtrl are based on Direct Inversion (Dir Inv) [17]. Our method achieves excellent performance in most metrics. Concretely, the results show that P2P has a better ability to preserve the contents of the source image but this also hinders the ability to edit images. As shown in CLIP scores, the P2P CLIP scores are relatively lower. In contrast, the PnP has a higher CLIP Edited score but the other metrics such as PSNR and LPIPS are worse. This indicates that PnP cannot preserve the invariance information. In summary, our method achieves good results in both background preservation and editing. This means that our method can make accurate editing without sacrificing either the invariance or the editing ability.

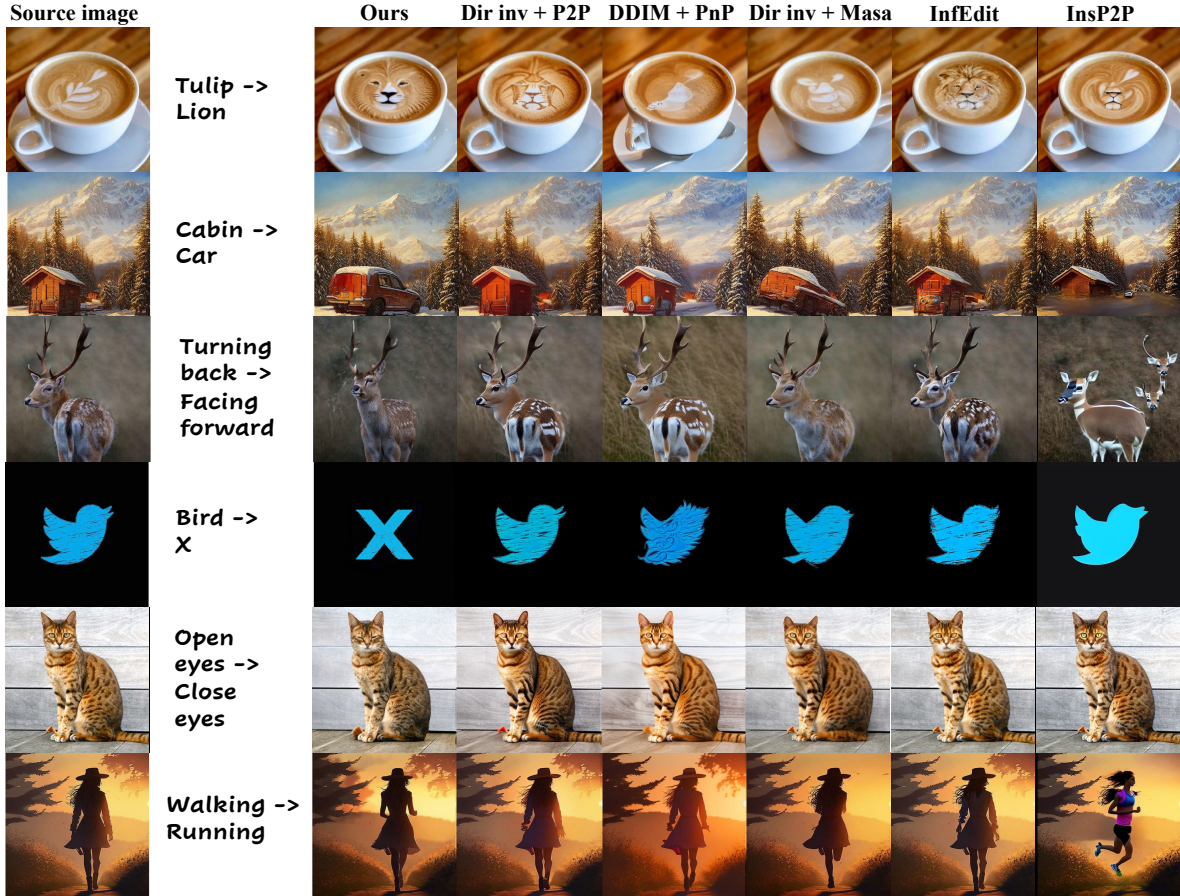


Figure 6. **Qualitative comparisons on real image editing.** With the proposed inversion and invariance control, our method can edit real images in different scenarios with different editing types, achieving high fidelity for different levels of geometric changes.

Non-rigid editing. Furthermore, to validate that our method can reconcile rigid and non-rigid editing scenarios, we specifically evaluate methods in non-rigid editing such as shape and pose with numerical experiments. We summarize the detailed results in the Appendix, and the comparison shows that our method outperforms others. This supports the motivation of our proposed invariance control module. The invariance control based on the AdaLN is aligned with the text changes so that the edited image can better adapt to non-rigid changes.

6.3. Ablation Study and Analysis

We ablate the proposed inversion and invariance control module to validate their effectiveness. We random sample 150 images from the PIE benchmark for ablation study.

Number of fixed-point iterations. We evaluate the effectiveness of the fixed-point iteration with different iteration numbers. The results are summarized in Table 2. Iter 0 means using the plain Euler inversion. The results show that the fixed-point iteration can improve the performance of preserving the unedited contents since the inversion becomes more accurate. Besides, a larger iteration number

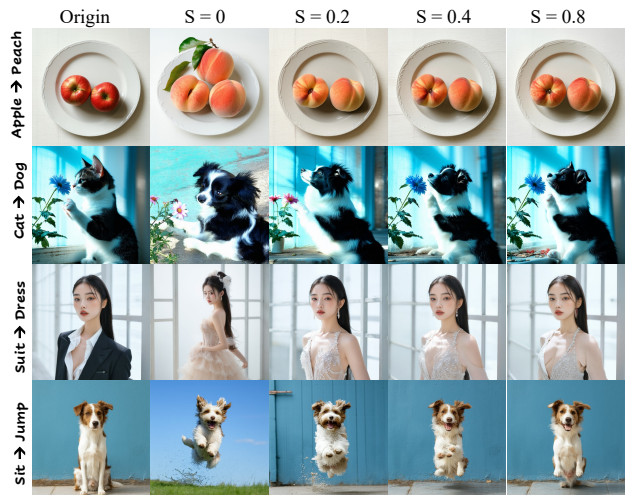


Figure 7. **Ablation study of invariance control of feature replacement in AdaLN under different time periods.**

can improve the performance, which is expected since more iterations improve the accuracy of inversion. However, we also observe that the benefit of using large numbers on in-

Table 2. **Ablation study on the fixed point iteration number and velocity compensation.** We choose different iteration numbers to validate their effectiveness in editing.

| Method | Structure | Background Preservation | | | | CLIP Similarity | |
|--------|------------|-------------------------|---------|--------|--------|-----------------|----------|
| | Distance ↓ | PSNR ↑ | LPIPS ↓ | MSE ↓ | SSIM ↑ | Whole ↑ | Edited ↑ |
| Iter 0 | 48.49 | 20.78 | 199.93 | 102.82 | 80.59 | 24.14 | 21.39 |
| Iter 1 | 15.29 | 26.98 | 67.80 | 35.68 | 93.89 | 24.31 | 21.47 |
| Iter 2 | 15.02 | 27.28 | 66.85 | 34.98 | 93.79 | 24.30 | 21.57 |
| Iter 3 | 14.92 | 27.51 | 66.48 | 34.06 | 94.40 | 24.48 | 21.73 |

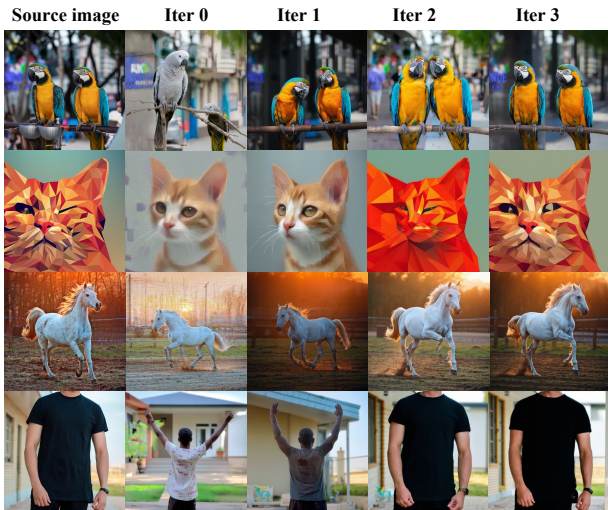


Figure 8. **Qualitative results of flow inversion under different iterations.** Fixed-point iteration can make the recovered image approximate the original image and then the left mismatch can be eliminated by the velocity compensation at each timestep.

variance preservation is marginal but this helps the editing since the CLIP scores are increased. We think the large number makes the inversion trajectory better converge to the authentic generation process and benefits editing. We also show the qualitative results of inversion under different iteration numbers in Figure 8. The results show that a large iteration number can facilitate the flow inversion. However, the larger iteration number also increases the computation load which needs to be considered for practical usage.

Velocity compensation. The velocity compensation eliminates the leftover error at each timestep when recovering the original image. The qualitative results in Figure 8 show that the left error is not large and the velocity compensation can recover the exact original image and improve the fidelity of editing. For the detailed qualitative results and analysis, please refer to Appendix.

Replacement of AdaLN within different time periods. To investigate the relation between invariance preservation with different time periods. We show the editing results with AdaLN invariance control within different time periods. The results in Figure 7 show that after a certain timestep, the edited image does not change a lot even if we

increase the time period. This indicates that the hyperparameter replacement timestep S is not sensitive. Besides, the results also show that applying text feature replacement within the AdaLN does not hinder non-rigid editing, which is better than the attention mechanism. So, we can set S to a relatively stable number in practice.

7. Limitations and Board Impact

Limitation. Tuning-free image editing is a system consisting of inversion and invariance control. One limitation is that the invariance control based on AdaLN depends on the authentic inversion resembling the generation process. If the real image is out of the model’s domain and the inversion deviates too much from the generation process, the module cannot precisely control the contents since the inversion does not fit the model’s prior distribution and text-to-image alignment is not good. Another limitation is the computational efficiency of the fixed-point iteration since this requires inferencing the transformer more times but in practice, the overhead is not large. We present the failure cases in the Appendix. We show the out-of-domain image with poor inversion. In this case, the model cannot well preserve the contents and manipulate the image as expected. The results are determined by the model’s prior, which is different from the original image.

Broad impact. Since the flow transformer has a large prior for the image in the real world, and our method can flexibly manipulate the images with text guidance, these raise concerns about the potential for creating misleading or deceptive content. This could be used maliciously in misinformation campaigns, deepfakes, or propaganda.

8. Conclusion

We investigated tuning-free image editing based on the flow transformer. We provide systematic analysis for the inversion of rectified flow and invariance modeling based on the MM-DiT. We propose a two-stage inversion strategy to reduce the approximation error of estimating the velocity field and try to find an inversion approximating the authentic generation process to enhance the editing ability. Besides, we propose the invariance control based on the AdaLN to reconcile rigid and non-rigid editing. The AdaLN control mechanism discriminately preserves the original image features according to the text changes, enabling more diverse editing types, especially for the rigid and non-rigid editing types. Both qualitative and quantitative experiments validate the effectiveness of the proposed modules. In conclusion, the proposed framework can adapt to versatile editing types and fully leverage the large generation prior of the flow transformer for tuning-free image editing. We hope our work can inspire thinking over flexible and accurate invariance control in large-scale T2I models to take full advantage of the generation prior for downstream tasks.

References

- [1] Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022.
- [2] 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. 2, 3, 5
- [3] Manuel Brack, Felix Friedrich, Katharia Kornmeier, Linoy Tsaban, Patrick Schramowski, Kristian Kersting, and Apolinário Passos. Ledits++: Limitless image editing using text-to-image models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8861–8870, 2024. 3, 4
- [4] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402, 2023. 1, 2, 3, 6
- [5] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22560–22570, 2023. 2, 3, 5, 6
- [6] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024. 1, 2, 3
- [7] Daniel Garibi, Or Patashnik, Andrey Voynov, Hadar Averbuch-Elor, and Daniel Cohen-Or. Renoise: Real image inversion through iterative noising. *arXiv preprint arXiv:2403.14602*, 2024. 3, 4
- [8] Zigang Geng, Binxin Yang, Tiankai Hang, Chen Li, Shuyang Gu, Ting Zhang, Jianmin Bao, Zheng Zhang, Houqiang Li, Han Hu, et al. Instructdiffusion: A generalist modeling interface for vision tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12709–12720, 2024. 1, 2, 3
- [9] Andrzej Granas, James Dugundji, et al. *Fixed point theory*. Springer, 2003. 4
- [10] Qin Guo and Tianwei Lin. Focus on your instruction: Fine-grained and multi-instruction image editing by attention modulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6986–6996, 2024. 2
- [11] 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. 1, 2, 3, 5, 6
- [12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 3
- [13] 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. 3
- [14] Vincent Tao Hu, Wei Zhang, Meng Tang, Pascal Mettes, Deli Zhao, and Cees Snoek. Latent space editing in transformer-based flow matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2247–2255, 2024. 2
- [15] Yuzhou Huang, Liangbin Xie, Xintao Wang, Ziyang Yuan, Xiaodong Cun, Yixiao Ge, Jiantao Zhou, Chao Dong, Rui Huang, Ruimao Zhang, et al. Smartedit: Exploring complex instruction-based image editing with multimodal large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8362–8371, 2024. 2
- [16] Inbar Huberman-Spiegelglas, Vladimir Kulikov, and Tomer Michaeli. An edit-friendly ddpm noise space: Inversion and manipulations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12469–12478, 2024. 3
- [17] 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. 3, 4, 6
- [18] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6007–6017, 2023. 1, 2
- [19] Gwanhyeong Koo, Sunjae Yoon, Ji Woo Hong, and Chang D Yoo. Flexiedit: Frequency-aware latent refinement for enhanced non-rigid editing. *arXiv preprint arXiv:2407.17850*, 2024. 2, 3, 5
- [20] Soyeong Kwon, Taegyeong Lee, and Taehwan Kim. Zero-shot text-guided infinite image synthesis with llm guidance. *arXiv preprint arXiv:2407.12642*, 2024. 2
- [21] Yuanze Lin, Yi-Wen Chen, Yi-Hsuan Tsai, Lu Jiang, and Ming-Hsuan Yang. Text-driven image editing via learnable regions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7059–7068, 2024. 2
- [22] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- [23] 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. 3
- [24] Jingyi Lu, Xinghui Li, and Kai Han. Regiondrag: Fast region-based image editing with diffusion models. *arXiv preprint arXiv:2407.18247*, 2024. 2
- [25] Nanye Ma, Mark Goldstein, Michael S Albergo, Nicholas M Boffi, Eric Vanden-Eijnden, and Saining Xie. Sit: Exploring flow and diffusion-based generative models with scalable interpolant transformers. *arXiv preprint arXiv:2401.08740*, 2024.
- [26] Zichong Meng, Changdi Yang, Jun Liu, Hao Tang, Pu Zhao, and Yanzhi Wang. Instructgie: Towards generalizable image editing. *arXiv preprint arXiv:2403.05018*, 2024. 2

- [27] 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
- [28] 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. 1, 3, 4
- [29] Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, and Ying Shan. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4296–4304, 2024.
- [30] 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. 1, 2
- [31] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023. 1
- [32] William S Peebles and Saining Xie. Scalable diffusion models with transformers. 2023 ieee. In *CVF International Conference on Computer Vision (ICCV)*, 2022.
- [33] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 1
- [34] Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman, and Ricky TQ Chen. Multisample flow matching: Straightening flows with minibatch couplings. *arXiv preprint arXiv:2304.14772*, 2023.
- [35] 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. 1
- [36] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015.
- [37] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding, 2022. *URL <https://arxiv.org/abs/2205.11487>*, 4. 1
- [38] Shelly Sheynin, Adam Polyak, Uriel Singer, Yuval Kirstain, Amit Zohar, Oron Ashual, Devi Parikh, and Yaniv Taigman. Emu edit: Precise image editing via recognition and generation tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8871–8879, 2024. 2
- [39] Yujun Shi, Chuhui Xue, Jun Hao Liew, Jiachun Pan, Han-shu Yan, Wenqing Zhang, Vincent YF Tan, and Song Bai. Dragdiffusion: Harnessing diffusion models for interactive point-based image editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8839–8849, 2024. 2
- [40] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2, 3
- [41] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 3
- [42] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1921–1930, 2023. 2, 3, 6
- [43] 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
- [44] 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. 3
- [45] Zongze Wu, Nicholas Kolkin, Jonathan Brandt, Richard Zhang, and Eli Shechtman. Turboedit: Instant text-based image editing. *arXiv preprint arXiv:2408.08332*, 2024. 3
- [46] Tianyu Xie, Yu Zhu, Longlin Yu, Tong Yang, Ziheng Cheng, Shiyue Zhang, Xiangyu Zhang, and Cheng Zhang. Reflected flow matching. *arXiv preprint arXiv:2405.16577*, 2024.
- [47] Sihan Xu, Yidong Huang, Jiayi Pan, Ziqiao Ma, and Joyce Chai. Inversion-free image editing with natural language. *arXiv preprint arXiv:2312.04965*, 2023. 2, 3, 5, 6
- [48] Sihan Xu, Yidong Huang, Jiayi Pan, Ziqiao Ma, and Joyce Chai. Inversion-free image editing with language-guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9452–9461, 2024. 2
- [49] Yifan Yang, Houwen Peng, Yifei Shen, Yuqing Yang, Han Hu, Lili Qiu, Hideki Koike, et al. Imagebrush: Learning visual in-context instructions for exemplar-based image manipulation. *Advances in Neural Information Processing Systems*, 36, 2024. 2
- [50] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*, 2023.
- [51] Kai Zhang, Lingbo Mo, Wenhui Chen, Huan Sun, and Yu Su. Magicbrush: A manually annotated dataset for instruction-guided image editing. *Advances in Neural Information Processing Systems*, 36, 2024. 1, 2