

High-Fidelity Diffusion-Based Image Editing

Chen Hou¹, Guoqiang Wei², Zhibo Chen^{1*}

¹University of Science and Technology of China

²ByteDance Research

houchen@mail.ustc.edu.cn, weiguoqiang.9@bytedance.com, chenzhibo@ustc.edu.cn

Abstract

Diffusion models have attained remarkable success in the domains of image generation and editing. It is widely recognized that employing larger inversion and denoising steps in diffusion model leads to improved image reconstruction quality. However, the editing performance of diffusion models tends to be no more satisfactory even with increasing denoising steps. The deficiency in editing could be attributed to the conditional Markovian property of the editing process, where errors accumulate throughout denoising steps. To tackle this challenge, we first propose an innovative framework where a rectifier module is incorporated to modulate diffusion model weights with residual features, thereby providing compensatory information to bridge the fidelity gap. Furthermore, we introduce a novel learning paradigm aimed at minimizing error propagation during the editing process, which trains the editing procedure in a manner similar to denoising score-matching. Extensive experiments demonstrate that our proposed framework and training strategy achieve high-fidelity reconstruction and editing results across various levels of denoising steps, meanwhile exhibits exceptional performance in terms of both quantitative metric and qualitative assessments. Moreover, we explore our model's generalization through several applications like image-to-image translation and out-of-domain image editing.

Introduction

As a rising star of generative models, tremendous works of diffusion models (Ho, Jain, and Abbeel 2020; Song et al. 2020) have been exploded in recent years. Except for works focusing on optimizing diffusion algorithm itself (Nichol and Dhariwal 2021; Song, Meng, and Ermon 2020), others devote to studying how to add controllable conditions to diffusion models, including adding image guidance (Choi et al. 2021), classifier guidance (Dhariwal and Nichol 2021; Avrahami, Lischinski, and Fried 2022), using representation learning (Kwon, Jeong, and Uh 2022) or additional networks (Rombach et al. 2022; Zhang and Agrawala 2023). These methods then inspire series of applications based on diffusion models, like image inpainting (Lugmayr et al. 2022), image translation (Meng et al. 2021), super resolution (Ho

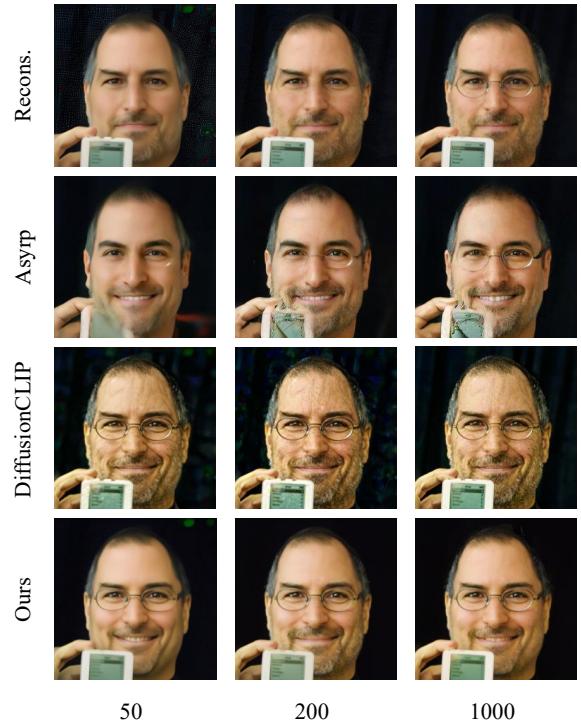


Figure 1: Reconstruction and editing results under various levels of inversion and denoising steps. While increasing steps makes reconstruction nearly perfect, the outcomes of editing still remain far from satisfactory (attribute: smiling).

et al. 2022) and image editing (Nichol et al. 2021; Kwon, Jeong, and Uh 2022; Hertz et al. 2022).

Existing work on image editing based on diffusion models could be roughly divided into two categories. One is through image guidance (Nichol et al. 2021; Yang et al. 2023), these methods take advantage of diffusion model's image-level noisy maps, and achieve editing by adding pixel-wise control through denoising process. But the disadvantages are that these methods need either mask (Avrahami, Lischinski, and Fried 2022), estimating mask (Couairon et al. 2022) or segmentation map (Matsunaga et al. 2022) to get fine control of images, besides, they are not suitable for semantic editing

*Corresponding author.

and the editing directions are usually heterogeneous. Other methods manipulate images via internal representation of diffusion, by exploring semantic latent space (Kwon, Jeong, and Uh 2022; Preechakul et al. 2022), or finetuning model parameters (Kim, Kwon, and Ye 2022; Kawar et al. 2023). These methods don't need mask as constraints, and except some of them only handle single image and corresponding text prompt as input (Hertz et al. 2022; Kawar et al. 2023), they get editing directions in good properties that are homogenous, linear and robust (Kwon, Jeong, and Uh 2022). Despite the superiorities, these methods, due to they offset the denoising process following editing directions, often cause changes in irrelevant attributes, and the details of image will also be lost or distorted. It should be noted that for reconstructions, increasing diffusion steps could do nearly perfect reconstruction for most images, but this doesn't hold true in the case of editing. The deficiency in editing could be attributed to its conditional Markovian property, leading to error accumulating and amplifying (Mokady et al. 2023). Fig.1 shows some reconstruction and editing results with various levels of diffuse and denoise steps from 50 (the lowest common steps used to save time) to 1000 (the highest steps adopted to train original DDPM), the editing attribute is smiling. As illustrated, reconstruction attains nearly perfect results with increasing steps, whereas for editing, the outcomes are still far from satisfactory.

To solve these problems, we first analyze why diffusion models suffer from distorted reconstructions or edits, and how these problems could be alleviated. Following that, we propose a designed framework and develop an effective training strategy to resolve these issues. Firstly we do this by adding a rectifier into diffusion model to fill the fidelity gap during denoising process. The rectifier is a hyper-network (David, Andrew, and Quoc 2016) that encodes the residual feature of original image and each step's estimation, at every step, it learns to predict the offsets of convolutional filters' weights for diffusion model's corresponding layers, providing compensated information for high-fidelity reconstruction. Secondly, to further reduce the propagation of error during editing process, we introduce a new paradigm for training editing based on diffusion models. Unlike previous methods who adopt Markov-like training strategies that make error accumulation (Kim, Kwon, and Ye 2022; Kwon, Jeong, and Uh 2022), we train editing in a way like denoising score matching (Song et al. 2020) which is widely used in training diffusion models (Ho, Jain, and Abbeel 2020; Song et al. 2020; Lipman et al. 2022). This restrains the trajectory deviation caused by editing not to accumulate, and effectively improves the faithfulness of edited results. Extensive experiments show that our method produces high-fidelity reconstruction and editing results without retraining diffusion model itself, especially for out-of-domain images.

To summarize, the main contributions are:

- We propose an innovative framework to achieve high-fidelity reconstruction and editing based on pretrained diffusion model, where a rectifier is incorporated to modulate model weights with residual features, providing compensated information for bridging the fidelity gap.

- To further reduce error propagation during editing, we propose a new learning paradigm where editing is trained in a manner similar to denoising score-matching. This prevents denoising trajectory from accumulated deviation, effectively improves the fidelity of edited results.

Related Work

Image Editing with Diffusion Models

The most intuitive way of using diffusion model for editing is to utilize the intermediate noisy maps generated during denoising process. These maps have the same resolutions as output images, making it convenient to directly add pixel-wise controls for manipulation, and their noisy property retains randomness for generation diversities. Many works take this advantage and apply it in various tasks like semantic editing (Choi et al. 2021), image translation (Meng et al. 2021), inpainting (Lugmayr et al. 2022), and pixel-level editing with mask (Nichol et al. 2021; Yang et al. 2023; Avrahami, Lischinski, and Fried 2022). Some other methods explore the influence of internal representation to attribute editing, instead of changing sampling process, they change the diffusion model itself by exploring the semantic latent inside (Kwon, Jeong, and Uh 2022), or finetuning the model to adapt editing tasks (Kim, Kwon, and Ye 2022; Hertz et al. 2022; Kawar et al. 2023). These methods could get homogenous and robust editing directions without the help of mask, but often suffer from distortion and low-fidelity. Besides interfering the denoising process or finetuning diffusion model, some methods take a novel yet different path to achieve editing by modulating the initial noise (Mao, Wang, and Aizawa 2023). There are also some novel methods who offer customized text control by inverting images into textual tokens (Gal et al. 2022a; Mokady et al. 2023).

High-Fidelity Inversion of GANs

Unlike the natural inversion capability exists in diffusion models (Song, Meng, and Ermon 2020), GANs (Goodfellow et al. 2020) need to do inversion by encoder (Richardson et al. 2021), optimization (Abdal, Qin, and Wonka 2020) or a combination of both (Zhu et al. 2020). Poor fidelity of inversion and reconstruction leads to the distortion-editability trade-off in GAN-based editing tasks (Tov et al. 2021). That is, good editing directions often lead to bad distortions and vice versa, it's hard to keep distortion and editing results both satisfying. Many works resolve this problem by improving inversion fidelity of GANs. Restyle (Alaluf, Patashnik, and Cohen-Or 2021) achieves this goal through iterative refine the residual of latent code. StyleRes (Pehlivan, Dalva, and Dundar 2023) transforms the residual of feature maps instead of images into editing branch, and propose a cycle-consistency loss to retain input details. HFGI (Wang et al. 2022) instead adaptively aligns the distortion map then fuses it into generator's internal feature maps, similar idea is also presented in ReGANIE (Li et al. 2023). While most works often keep generator weight unchanged, there are other methods like HyperStyle (Alaluf et al. 2022) who chooses to finetune generator parameters. Motivated by

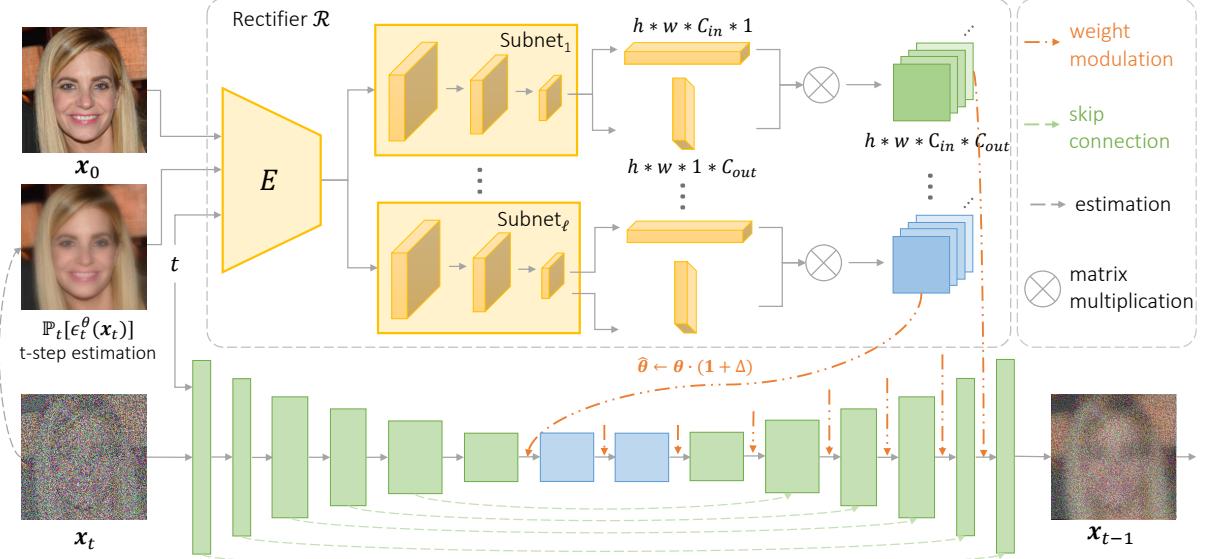


Figure 2: Overview of our proposed rectifier framework. The rectifier is a hypernetwork consisting of a global encoder and multiple subnet branches. It takes as input the original image x_0 and the estimation at each step ($\mathbb{P}_t[\epsilon_t^\theta(x_t)]$), targets to modulate the degraded residual features into offset weights, providing compensated information for high-fidelity reconstruction. We select the middle and up-sampling blocks of U-Net for modulate, considering that these blocks contain both high-level semantic information and low-level details. We also employ separable convolution to reduce the amount of generated parameters.

these methods, while considering the particularity of diffusion model relative to GAN, we propose a high-fidelity framework adapted for diffusion models.

Methodology

In this section, we start by explaining why diffusion models suffer from distorted reconstructions and edits. Then we will elaborate on how these issues could be alleviated, followed by the introduction of our method.

High-Fidelity Problem in Diffusion Models

Reconstructions of diffusion model are not always perfect. As claimed in PDAE (Zhang, Zhao, and Lin 2022), the major reason of these imperfections is there exists a clear gap between the predicted posterior mean and the true one. Compared to reconstruction, editing deviates denoising trajectory thus leads to more error accumulation (Mokady et al. 2023), while (Ho and Salimans 2022) also figures out that the effect induced by editing condition (such as classifier or conditioned text) will be amplified during denoising process, making editing a harder task than merely reconstruction.

According to (Zhang, Zhao, and Lin 2022), some prior knowledge about x_0 introduced to the reverse process will help reduce the gap and achieve better reconstruction. From this perspective, classifier-guidance (Dhariwal and Nichol 2021) method can be seen as utilizing the class information to fill this gap, via shifting the predicted posterior mean with an extra item computed by the classifier’s gradient (Zhang, Zhao, and Lin 2022). PDAE also proves that this is equivalent to shifting the noise predicted by the model, thus they

use an additional network predicting noise shift to make up for the information loss.

Hypernetwork as Rectifier

While these methods point out how to compensate for the information gap, their reconstructions are still far from high-fidelity, and resorting to external network or classifier also makes it difficult to generalize to semantic editing tasks which mainly relies on diffusion’s internal representations. In this work, we propose a framework where a rectifier is incorporated to modulate residual features into offset weights, providing compensated information to help pretrained diffusion model achieving high-fidelity reconstructions. The framework is illustrated in Fig. 2. Our rectifier is a hypernetwork (David, Andrew, and Quoc 2016) which takes as input every step’s estimation and original image, expected to exploit the degraded residual features to fill the fidelity gap. The inputs first pass through a global encoder, then transformed by a series of sub-nets to generate layer-wise modulation. We choose to modulate the middle and up-sampling blocks of U-Net (Ronneberger, Fischer, and Brox 2015), considering they contain both high-level semantic information and low-level details. Furthermore, without the interference of other representations like classifier, our framework is highly adaptive for semantic editing tasks, and is easier to generalize to other diffusion-based downstream tasks.

For parameter modulation, we generate offsets for all convolutional layers’ kernel weights, instead of regenerating them from scratch. This can preserve prior knowledge of pretrained diffusion model as much as possible (Alaluf et al. 2022). Specifically, at time step t , the rectifier \mathcal{R} takes in the

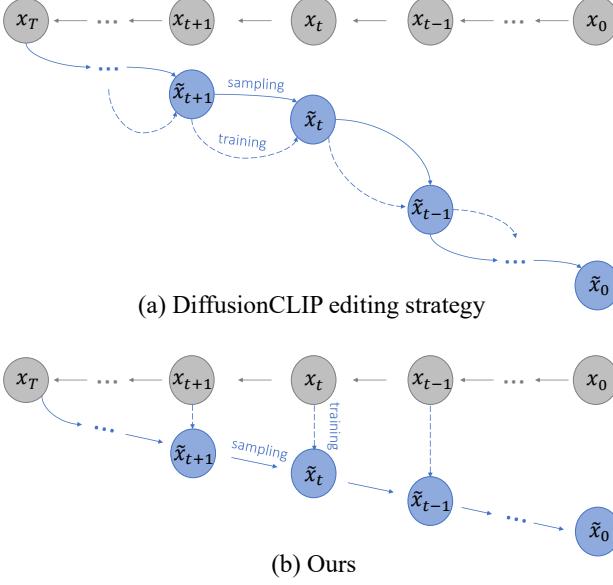


Figure 3: Editing training strategy. Instead of shifting from previous edited results in a Markovian style used in DiffusionCLIP (a), which may lead to error propagation, we start from the original trajectory at each step to find editing direction (b), further alleviating error accumulation caused in editing process.

original image x_0 and the prediction result using x_t , then outputs weight offsets Δ_t for ℓ -th layer of U-Net which are then assigned to each channel i of j -th filter:

$$\Delta_{\ell,t}^{i,j} := \mathcal{R}(x_0, \mathbb{P}_t[\epsilon_t^\theta(x_t)], t), \quad (1)$$

where ϵ_t^θ represents the noise estimation at time step t with parameter θ . $\mathbb{P}_t[\epsilon_t^\theta(x_t)] = (x_t - \sqrt{1 - \alpha_t} \epsilon_t^\theta) / \sqrt{\alpha_t}$ refers to the estimation of x_0 using x_t defined in DDIM (Song, Meng, and Ermon 2020), and α_t denotes the transformed variance noise schedule used in DDPM (Ho, Jain, and Abbeel 2020). The kernel weight is modulated as:

$$\theta_{\ell,t}^{i,j} := \theta_{\ell,t} \cdot (1 + \Delta_{\ell,t}^{i,j}). \quad (2)$$

Considering the huge cost of estimating weight offsets for all selected layers, we employ separable convolution (Alaluf et al. 2022) to cut down the amount of parameters generated. Rather than predicting offsets for every filter of every channel (which requires $\sum_\ell h * w * C_{in} * C_{out}$ parameters generated in total), we decompose it into two parts: $h * w * C_{in} * 1$ and $h * w * 1 * C_{out}$, their product is taken as the final output. In this way, the number of parameters is reduced to $\sum_\ell (h * w * C_{in} * 1 + h * w * 1 * C_{out})$. This significantly reduces memory usage of the network, while not affecting its capability too much. For loss function, we choose noise fitting loss as our training objective:

$$\mathcal{L}_{rec} := \mathbb{E}_{t,x_0,\epsilon} \left[\left\| \epsilon - \epsilon_t^\theta(x_t) \right\|_2^2 \right]. \quad (3)$$

It is rational to consider other loss functions, like ℓ_1 loss, which is commonly used in GAN finetuning tasks (Alaluf

Algorithm 1: Editing Training Strategy

```

1: repeat
2:    $x_0 \sim q(x_0)$ 
3:    $t \sim \text{Uniform}(\{1, \dots, T\})$ 
4:    $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
5:    $\tilde{x}_t = \sqrt{\alpha_t} x_0 + \sqrt{1 - \alpha_t} \epsilon$ 
6:    $\tilde{\theta} \leftarrow \theta \cdot (1 + \mathcal{R}(x_0, \mathbb{P}_t[\epsilon_t^\theta(x_t)], t))$ 
7:   Take gradient descent step on
       $\nabla_{\mathcal{R}} \mathcal{L}_{\text{direction}}(\mathbb{P}_t[\epsilon_t^\theta(x_t)], t_{tar}; x_0, t_{src})$ 
       $\nabla_{\mathcal{R}} \mathcal{L}_{\ell_1}(\mathbb{P}_t[\epsilon_t^\theta(x_t)], x_0)$ 
8: until converged

```

et al. 2022; Wang et al. 2022). We also validate the effectiveness of these loss functions for diffusion models, and the relevant results are shown in supplementary material.

Training Editing Like Score Matching

As claimed before, compared with reconstruction, editing is more challenging and more susceptible to causing distortion during denoising process, owing to the error accumulation introduced by input condition. How to diminish these impact and keep high-fidelity for editing is a critical issue we should consider next. Current methods train editing in a Markovian way (Kim, Kwon, and Ye 2022; Kwon, Jeong, and Uh 2022), in which case the deviation of denoising trajectory will gradually accumulate, leading to irrelevant attributes change, details loss or distortion (Fig.1). To alleviate this problem and further reduce the error propagation in editing process, we propose a training strategy that trains editing in a manner similar to denoising score matching (Song et al. 2020). Our editing training strategy is depicted in Fig.3. The inspiration is drawn from the training strategy of diffusion model (Ho, Jain, and Abbeel 2020) and score-based generative model (Song et al. 2020). Instead of drifting from previous edited results in a Markovian way, we instead take the original trajectory as the starting point to find editing directions for each step. This eschews the accumulation of the the deviations caused by editing, and further reduces the error propagated in editing process. Specifically, we reuse the rectifier to modulate model’s weights served for editing, which can also be interpreted as shifting the output distribution of the entire model along the direction of attribute change.

Another advantage of our editing training strategy is that we do not need to specify any heuristic-defined parameters to fit different attributes. It should be remarked here that for methods like Asymp (Kwon, Jeong, and Uh 2022) and DiffusionCLIP (Kim, Kwon, and Ye 2022), neither of them employs editing through the entire denoising process. Asymp halts editing prematurely and adds stochastic noise then to boost perceived quality, while DiffusionCLIP does not inverse images into complete noise for preserving their morphologies. Setting these parameters meticulously for every separate attribute is intricate and bothersome. Our method though, starts from pure noise and traverses the process thoroughly to get editing results, no extra settings are needed.

We incorporate the directional CLIP loss (Gal et al.

2022b) to train the editing process. Specifically, given the source image \mathbf{x}_{src} and text t_{src} as well as the target image \mathbf{x}_{tar} and text t_{tar} , we can calculate the feature directions encoded by CLIP’s image encoder E_I and text encoder E_T , *i.e.*, $\Delta I = E_I(\mathbf{x}_{tar}) - E_I(\mathbf{x}_{src})$ and $\Delta T = E_T(t_{tar}) - E_T(t_{src})$. The directional CLIP loss aims to align the image change ΔI and text change ΔT , which could be formulated as:

$$\mathcal{L}_{direction}(\mathbf{x}_{tar}, t_{tar}; \mathbf{x}_{src}, t_{src}) := 1 - \frac{\langle \Delta I, \Delta T \rangle}{\|\Delta I\| \|\Delta T\|}. \quad (4)$$

Motivated by (Kim, Kwon, and Ye 2022), we also introduce another ℓ_1 loss as a regularizer to circumvent the change in irrelevant attributes:

$$\mathcal{L}_{\ell_1}(\mathbf{x}_{tar}, \mathbf{x}_{src}) := \|\mathbf{x}_{tar} - \mathbf{x}_{src}\|. \quad (5)$$

Our final loss function for training editing is:

$$\mathcal{L}_{edit} := \lambda_{CLIP} \mathcal{L}_{direction} + \lambda_{recon} \mathcal{L}_{\ell_1}. \quad (6)$$

Training of editing is established upon the foundation of model pretrained by the rectifier part. During inference, we still use the same sampling procedure as DDPM (Ho, Jain, and Abbeel 2020), but with the modulated model that leads to corresponding attribute change. Our training strategy is elucidated in Algorithm 1.

Experiments

Implementation Details

We conduct experiments on FFHQ (Karras, Laine, and Aila 2019), CelebA-HQ (Karras et al. 2017), AFHQ-dog (Choi et al. 2020), METFACES (Karras et al. 2020), LSUN-church/-bedroom (Yu et al. 2015) datasets with the outcomes of various levels of steps, and all pretrained models are kept frozen. Note that due to the separable convolution used in rectifier, our model is GPU-efficient and are able to complete all training tasks on a single RTX 3090TI GPU.

Reconstructions

We present both quantitative and qualitative evaluations of image reconstruction. We conduct our rectifier on several backbones with various datasets, and the quantitative results are shown in Table 1. iDDPM (Nichol and Dhariwal 2021) is employed for human faces, and the metrics are calculated on 10,000 random sampled images from CelebA-HQ using model trained on FFHQ under 50 inversion and sampling steps. In terms of natural scenes, we use DDPM++ (Song et al. 2020) as foundation model and implement on LSUN-Church (Yu et al. 2015) with 20 steps. It is worth noting that even though we do not train on these indicators and only train with noise fitting loss as Eq.(1), our method still outperforms original model under some of the reconstruction quality assessment criterias. We also test the average posterior mean gap $\|\Delta e\|^2$, and it turns out our method reaches lower gap than original model. These results manifest our rectifier could bring quality improvement for model’s overall output distribution, and indeed provides compensated information thus fills the fidelity gap.

Method	L_1	L_2	LPIPS	SSIM	$\ \Delta e\ ^2$
iDDPM	0.090	0.016	0.150	0.95	6.671e-3
Ours	0.085	0.014	0.150	0.94	6.671e-3
DDPM++	0.255	0.109	0.643	0.48	1.559e-2
Ours	0.254	0.108	0.642	0.48	1.558e-2

Table 1: Quantitative results of image reconstruction.



Figure 4: Comparison of reconstruction quality under 50 steps. Our method is more robust to occlusions (1st column), illuminations (2nd column), viewpoints (3rd and 4th columns), and performs better at restoring coarse shapes (5th column) and preserving fine details (6th column).

Some qualitative samples are shown in Fig.4. With the help of rectifier, our reconstructions become robust to occlusions, illuminations, viewpoints, and performs better at both restoring coarse shapes and preserving details. More visual results can be found in the supplementary materials.

Editings

For comparison of editing performance, we choose the representational editing methods based on diffusion backbones that retain state-of-the-art: Asyrp (Kwon, Jeong, and Uh 2022), DiffusionCLIP (Kim, Kwon, and Ye 2022). Among them, Asyrp leverages the deepest feature maps inside UNet’s bottleneck, treating it as diffusion model’s semantic latent space to produce manipulations. DiffusionCLIP, on the other hand, directly finetunes the whole model for attaining editing results. We also conduct experiment on some image-guidance methods like GLIDE (Nichol et al. 2021) to test their abilities towards semantic editing, for which the results could be found in supplementary materials.

In a comparable manner to reconstruction part, both quantitative and qualitative evaluation of editing are exhibited here too. Fig.5 presents some qualitative comparisons towards different methods trained under 50 inversion and sampling steps. Like previously noted, neither of Asyrp nor DiffusionCLIP employs edit through the entire denoising process, Asyrp applies stochastic noise during final process to boost perceived quality, and DiffusionCLIP begins editing from intermediate noisy images for preserving their morphologies. Our method, though starts from complete noise and traverses the whole denoising process to edit, still carries out editing results with exceptional quality. Several illustrative examples are, for instance, elements like hat and ring are kept intact during semantic editing, along with the preserva-

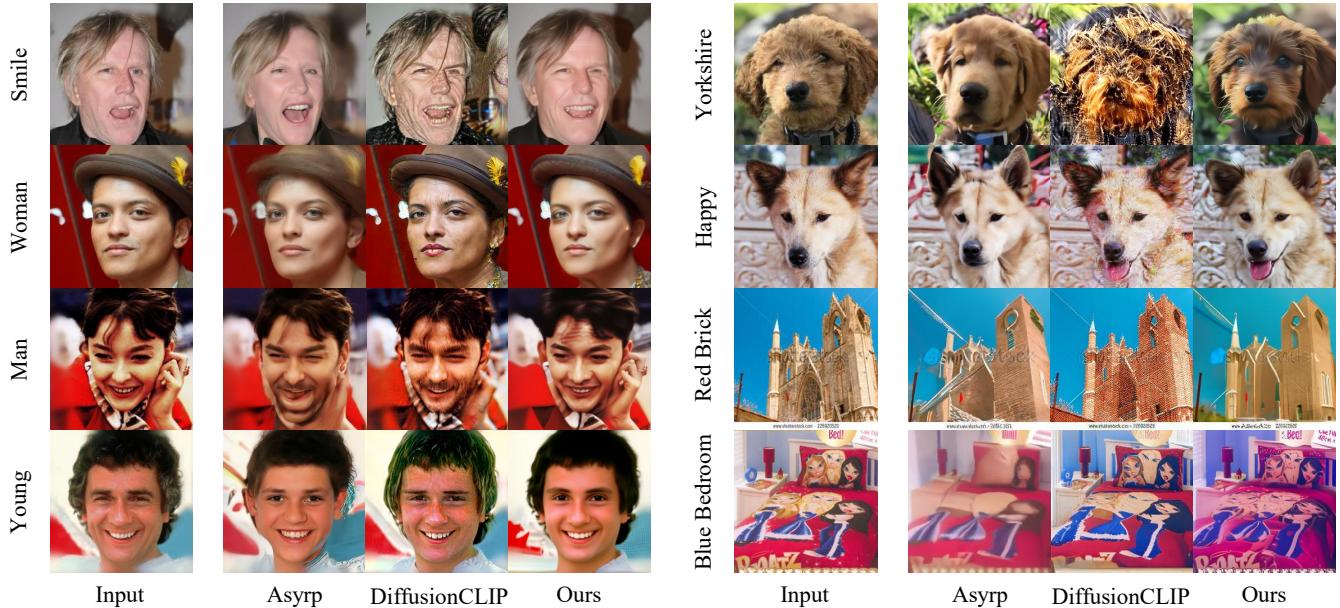


Figure 5: Editing qualitative comparisons. Our method delivers realistic edits while maintaining low distortion and high fidelity.

Attribute	Asyrr	DiffusionCLIP	Ours
Man	0.22 ± 0.16	0.33 ± 0.14	0.45 ± 0.23
Pixar	0.18 ± 0.13	0.22 ± 0.13	0.25 ± 0.10

Table 2: Quantitative comparisons of editing. We compare different methods with the identity similarity between original and edited images.

tion of image's overall shape and background. Meanwhile, distortions or artifacts brought by conditional input text are avoided, and vital information loss is also alleviated.

Echoing what's previously mentioned, editing as a more challenging task compared to reconstruction, its quality does not tend to improve much even with increasing inversion and denoising steps, mainly due to the error accumulation introduced by input conditions. In order to reinforce this standpoint, we test the editing performance under various levels of inversion and sampling steps from 50 to 1000. The outcomes are highlighted in Fig.1. As can be observed, methods like Asyrr and DiffusionCLIP who use Markovian training strategy fail to produce realistic and faithful editing results, even with larger steps. Asyrr loses many essential information like the iPod in hand and the glasses. DiffusionCLIP benefits some details from its incomplete noise inversion, yet still leads to distortions and artifacts. Our method attains vivid editing results regardless the number of steps, meanwhile maintaining high-fidelity performance in preserving vital information and details.

We offer quantitative results as well. Given original images and their edited versions, we calculate the identity similarity using CurricularFace (Huang et al. 2020), which grants us the capability to validate how identity are preserved before and after editing. Two attributes: man and



Figure 6: Effects of different editing training strategies. Different methods are evaluated across various ranges of editing intervals, "original" denotes default configuration, while "full" refers to editing through the entire denoising process. View with better clarity when zoomed-in.

pixar are evaluated, and all other methods are tested under official checkpoints. Table 2 showcases the results. It is evident from the table that out method obtains the highest identity similarity score among these attributes. Random sampled images used to calculate identity similarity and more details are shown in supplementary materials.

Ablation Study

Effect of Editing Training Strategy Doubts may arise regarding to whether our gain in editing comes from the rectifier or the training strategy. With the benefit of rectifier already proven in preceding part, we now focus on ablation studies to validate the effectiveness of editing training strategy alone. Attribute "smiling" is chosen for evaluation of various methods, and Fig.6 presents the outputs of these algorithms. As shown, even without any enhancement from rectifier, our strategy still produces results with less distor-

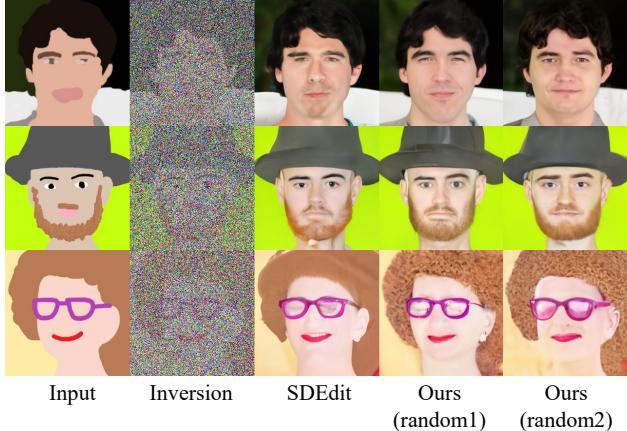


Figure 7: The influence of incorporating rectifier into SDEdit. The rectifier makes translation results more lifelike and realistic, as well as exhibiting richer texture and details. No extra domain specific training are employed.

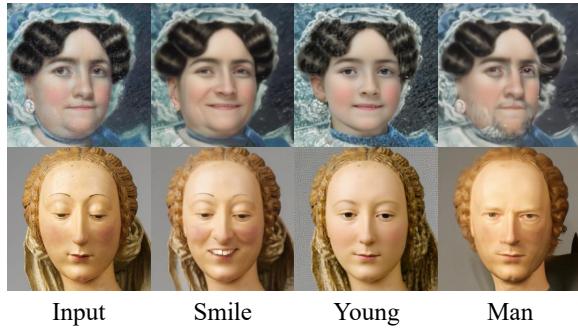


Figure 8: Generalize our method to out-of-domain images. Our model trained only on FFHQ successfully adapts to images from METFACES, performing well on oil paintings and sculptures which possess intricate and unique textures unseen in FFHQ.

tion and information loss compared to others. This demonstrates our strategy could be employed as an independent and generalized training approach for editing tasks based on diffusion models. Furthermore, recall that neither Asymp nor DiffusionCLIP implements editing through the entire process, they either stops prematurely or starts from incomplete noise. We thus investigate how they perform when applied to full range editing, denoted as "full" in Fig.6. It turns out that longer editing interval does not yield satisfying editing results. Especially for method like DiffusionCLIP, extending interval instead leads to loss of details and many artifacts. This observation proves again that improving performance of editing based on diffusion model necessitates more than simply increasing the editing interval steps.

Further Applications

Image to Image Translation The rectifier module can be incorporated into any pretrained diffusion models to en-

hance their quality for overall output distribution, indicating its potential for generalizing to various downstream tasks that utilize diffusion model as basis. One of these tasks involves images translation. SDEdit (Meng et al. 2021) firstly exploits the advantage of diffusion's stochasticity and the prior knowledge hidden in pretrained model, making translation task simple to achieve. Here, we perform image translation in the same way SDEdit does, but with rectifier integrated, in order to evaluate its capability generalizing to other tasks. Noted that no additional domain-specific training are employed in this scenario, and we only adopt the rectifier pretrained on FFHQ dataset. The results are shown in Fig.7. Benefiting from the rectifier, the translation results become more realistic and exhibit richer in texture and details (like the hat and the hair). This inspiring finding demonstrates that our rectifier module indeed learns to produce compensated information, and possesses the capability of generalizing to other downstream tasks, bringing further quality enhancement for them.

Generalization on Out-of-Domain Images For a more extensive evaluation of how our method generalizes, we further test its performance on images from other domains. Here we select images from METFACES and use our method pretrained on FFHQ dataset to edit. These out-of-domain images including oil paintings with complicated texture and details, as well as sculptures that possesses unique tactile qualities. We find out that even without any adjustment or finetuning for the new domain, our model could give expected outcomes achieving dual advantages in both editing performance and fidelity preservation. As depicted in Fig.8, while obtaining realistic and faithful edits, our method preserves greatly the intricate details such as texture of clothing, style of hair, together with images' whole structures. This signifies our method could handle diverse images from various similar domains, without explicitly finetuning on it, demonstrating its strong generalization ability.

Conclusions

In this work, we propose an innovative method to achieve high-fidelity image reconstruction and editing based on diffusion models. We employ a rectifier to encode residual feature into modulated weight, bringing compensated information for filling the fidelity gap. Furthermore, we introduce an effective editing learning paradigm which trains editing in a way like denoising score-matching, preventing error accumulation during editing process. By leveraging the rectifier and the training paradigm, our method produces high-fidelity reconstruction and editing results regardless of inversion and sampling steps. Comprehensive experiments validates the effectiveness of our method, and shows its strong generalization ability for editing out-of-domain images, or improving quality for various downstream tasks based on diffusion models.

Acknowledgements

This work was supported partly by Natural Science Foundation of China (NSFC) under Grant 62371434, 62021001.

References

Abdal, R.; Qin, Y.; and Wonka, P. 2020. Image2stylegan++: How to edit the embedded images? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8296–8305.

Alaluf, Y.; Patashnik, O.; and Cohen-Or, D. 2021. Restyle: A residual-based stylegan encoder via iterative refinement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6711–6720.

Alaluf, Y.; Tov, O.; Mokady, R.; Gal, R.; and Bermano, A. 2022. Hyperstyle: Stylegan inversion with hypernetworks for real image editing. In *Proceedings of the IEEE/CVF conference on computer Vision and pattern recognition*, 18511–18521.

Avrahami, O.; Lischinski, D.; and Fried, O. 2022. Blended diffusion for text-driven editing of natural images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18208–18218.

Choi, J.; Kim, S.; Jeong, Y.; Gwon, Y.; and Yoon, S. 2021. Ilvr: Conditioning method for denoising diffusion probabilistic models. *arXiv preprint arXiv:2108.02938*.

Choi, Y.; Uh, Y.; Yoo, J.; and Ha, J.-W. 2020. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8188–8197.

Couairon, G.; Verbeek, J.; Schwenk, H.; and Cord, M. 2022. Diffedit: Diffusion-based semantic image editing with mask guidance. *arXiv preprint arXiv:2210.11427*.

David, H.; Andrew, D.; and Quoc, V. 2016. Hypernetworks. *arXiv preprint arXiv:1609*.

Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.

Gal, R.; Alaluf, Y.; Atzmon, Y.; Patashnik, O.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2022a. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*.

Gal, R.; Patashnik, O.; Maron, H.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2022b. StyleGAN-NADA: CLIP-guided domain adaptation of image generators. *ACM Transactions on Graphics (TOG)*, 41(4): 1–13.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.

Hertz, A.; Mokady, R.; Tenenbaum, J.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2022. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*.

Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.

Ho, J.; Saharia, C.; Chan, W.; Fleet, D. J.; Norouzi, M.; and Salimans, T. 2022. Cascaded diffusion models for high fidelity image generation. *The Journal of Machine Learning Research*, 23(1): 2249–2281.

Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.

Huang, Y.; Wang, Y.; Tai, Y.; Liu, X.; Shen, P.; Li, S.; Li, J.; and Huang, F. 2020. Curricularface: adaptive curriculum learning loss for deep face recognition. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5901–5910.

Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2017. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.

Karras, T.; Aittala, M.; Hellsten, J.; Laine, S.; Lehtinen, J.; and Aila, T. 2020. Training generative adversarial networks with limited data. *Advances in neural information processing systems*, 33: 12104–12114.

Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4401–4410.

Kawar, B.; Zada, S.; Lang, O.; Tov, O.; Chang, H.; Dekel, T.; Mosseri, I.; and Irani, M. 2023. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6007–6017.

Kim, G.; Kwon, T.; and Ye, J. C. 2022. Diffusionclip: Text-guided diffusion models for robust image manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2426–2435.

Kwon, M.; Jeong, J.; and Uh, Y. 2022. Diffusion models already have a semantic latent space. *arXiv preprint arXiv:2210.10960*.

Li, B.; Ma, T.; Zhang, P.; Hua, M.; Liu, W.; He, Q.; and Yi, Z. 2023. ReGANIE: Rectifying GAN Inversion Errors for Accurate Real Image Editing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 1269–1277.

Lipman, Y.; Chen, R. T.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2022. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*.

Lugmayr, A.; Danelljan, M.; Romero, A.; Yu, F.; Timofte, R.; and Van Gool, L. 2022. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11461–11471.

Mao, J.; Wang, X.; and Aizawa, K. 2023. Guided Image Synthesis via Initial Image Editing in Diffusion Model. *arXiv preprint arXiv:2305.03382*.

Matsunaga, N.; Ishii, M.; Hayakawa, A.; Suzuki, K.; and Narihira, T. 2022. Fine-grained Image Editing by Pixelwise Guidance Using Diffusion Models. *arXiv preprint arXiv:2212.02024*.

Meng, C.; He, Y.; Song, Y.; Song, J.; Wu, J.; Zhu, J.-Y.; and Ermon, S. 2021. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*.

Mokady, R.; Hertz, A.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2023. Null-text inversion for editing real images using guided diffusion models. In *Proceedings of the*

IEEE/CVF Conference on Computer Vision and Pattern Recognition, 6038–6047.

Nichol, A.; Dhariwal, P.; Ramesh, A.; Shyam, P.; Mishkin, P.; McGrew, B.; Sutskever, I.; and Chen, M. 2021. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*.

Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, 8162–8171. PMLR.

Pehlivan, H.; Dalva, Y.; and Dundar, A. 2023. StyleRes: Transforming the residuals for real image editing with stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1828–1837.

Preechakul, K.; Chatthee, N.; Wizadwongsa, S.; and Suwanjanakorn, S. 2022. Diffusion autoencoders: Toward a meaningful and decodable representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10619–10629.

Richardson, E.; Alaluf, Y.; Patashnik, O.; Nitzan, Y.; Azar, Y.; Shapiro, S.; and Cohen-Or, D. 2021. Encoding in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2287–2296.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.

Ronneberger, O.; Fischer, P.; and Brox, T. 2015. 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, 234–241. Springer.

Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.

Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.

Tov, O.; Alaluf, Y.; Nitzan, Y.; Patashnik, O.; and Cohen-Or, D. 2021. Designing an encoder for stylegan image manipulation. *ACM Transactions on Graphics (TOG)*, 40(4): 1–14.

Wang, T.; Zhang, Y.; Fan, Y.; Wang, J.; and Chen, Q. 2022. High-fidelity gan inversion for image attribute editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11379–11388.

Yang, B.; Gu, S.; Zhang, B.; Zhang, T.; Chen, X.; Sun, X.; Chen, D.; and Wen, F. 2023. Paint by example: Exemplar-based image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18381–18391.

Yu, F.; Seff, A.; Zhang, Y.; Song, S.; Funkhouser, T.; and Xiao, J. 2015. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*.

Zhang, L.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. *arXiv preprint arXiv:2302.05543*.

Zhang, Z.; Zhao, Z.; and Lin, Z. 2022. Unsupervised representation learning from pre-trained diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 35: 22117–22130.

Zhu, J.; Shen, Y.; Zhao, D.; and Zhou, B. 2020. In-domain gan inversion for real image editing. In *European conference on computer vision*, 592–608. Springer.