

Not All Steps are Created Equal: Selective Diffusion Distillation for Image Manipulation

Luozhou Wang^{1, 4*} Shuai Yang^{1, 4*} Shu Liu² Ying-cong Chen^{1, 3, 4†}

¹ The Hong Kong University of Science and Technology (Guangzhou). ² SmartMore.

³ The Hong Kong University of Science and Technology. ⁴ HKUST (Guangzhou) - SmartMore Joint Lab.

Abstract

Conditional diffusion models have demonstrated impressive performance in image manipulation tasks. The general pipeline involves adding noise to the image and then denoising it. However, this method faces a trade-off problem: adding too much noise affects the fidelity of the image while adding too little affects its editability. This largely limits their practical applicability. In this paper, we propose a novel framework, *Selective Diffusion Distillation (SDD)*, that ensures both the fidelity and editability of images. Instead of directly editing images with a diffusion model, we train a feedforward image manipulation network under the guidance of the diffusion model. Besides, we propose an effective indicator to select the semantic-related timestep to obtain the correct semantic guidance from the diffusion model. This approach successfully avoids the dilemma caused by the diffusion process. Our extensive experiments demonstrate the advantages of our framework. Code is released at <https://github.com/AndysonYs/Selective-Diffusion-Distillation>.

1. Introduction

In recent years, diffusion model [10, 14, 28, 34, 36, 38, 40–43] has attracted great attention in both academic and industrial communities. It models the Markov transition from a Gaussian distribution to a data distribution to generate high-quality images sequentially. The elegant formulation achieves state-of-the-art performance in various image generation benchmarks. Meanwhile, text-to-image diffusion models [28, 34, 36, 38] also demonstrate their impressive capacity in controllable image synthesis, enabling a wide range of practical applications. Among them, one of the most interesting applications is image manipulation.

A typical text-guidance manipulation pipeline is to invert the input image to a noisy latent and then denoise this latent

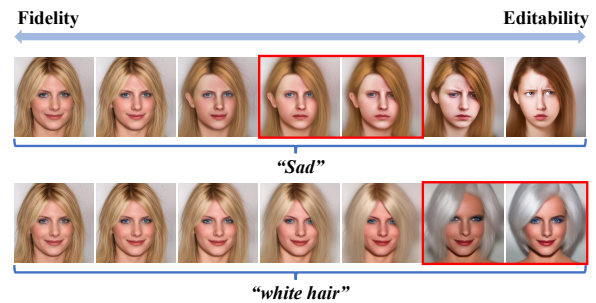


Figure 1. **Editability and fidelity trade-off of diffusion-based image manipulation.** The leftmost is the input image. For each manipulation, we add increasing noise levels from left to right and then denoise the image. Different semantics require different levels of noise to manipulate.

with a given text prompt. The inversion process can be the simple noise adding [25] or the DDIM inversion [41, 44]. There is an editability-fidelity tradeoff in such pipelines, as shown in Fig. 1. Adding a lot of noise to the original image gives the diffusion model more freedom to manipulate the image, but it also makes it harder to retain original semantics when denoising back, and vice versa.

More importantly, this trade-off can lead to unsuccessful manipulations, such as “white hair” in Fig. 1. This is because the diffusion model processes different semantics at different stages [6]. If our manipulation corresponds to the semantics in the early stages, we have to add more noise, hence losing much information from the original image. To overcome the failure caused by the trade-off, existing methods attempt to incorporate more guidance information, such as using masks to limit the manipulation region [1, 2, 24]. This is useful for local image manipulation, but for some global structures, such as changing the pose of a human face, it still fails to solve the problem.

In this work, we take a different viewpoint to leverage a diffusion model for image manipulation. Unlike existing methods that directly manipulate images progressively with diffusion models, we train an efficient image manipulator

*Equal contribution

†Corresponding author

supervised by a pretrained diffusion model. Specifically, our model takes a feedforward model (e.g., latent space image manipulation model [29]) as the manipulator, and a text-guided diffusion model (e.g., latent diffusion [36]) as the supervisor. During training, the manipulated image is diffused and fed to the diffusion model, and the diffusion model produces gradient supervision based on the text condition. In this sense, it is expected that the manipulator could mimic the promising generation capacity of the diffusion model.

To learn such a manipulator, using correct semantic guidance is also crucial, as shown in Fig. 1. Intuitively, the diffusion model at the timestep in the red box has a better ability to guide the image manipulator in changing the hair color. In contrast, other timesteps do not provide useful semantic guidance. Therefore, we propose the hybrid quality score (HQS), an effective indicator that helps to select appropriate timesteps. This indicator is built upon the entropy and L_1 norm of the gradient from the diffusion model and is shown to be highly correlated with the manipulation quality. As such, our model could learn with the most effective guidance. In addition to solving the trade-off problem of diffusion-based image manipulation, we have another bonus: our image manipulator requires only one forward for manipulation, and by learning in a specific domain, our image manipulator can manipulate images over the entire domain, which offsets the cost of training. Extensive experiments demonstrate the effectiveness of our methods.

Our contributions are summarized as follows:

- We propose a novel image manipulation approach with a well-trained diffusion model to supervise another image manipulator. This avoids the trade-off problem of manipulating images with a diffusion process.
- We propose the hybrid quality score to detect semantic-related timesteps. Only during these timesteps can the diffusion model guide the image manipulator to perform accurate manipulations.
- Our experiments demonstrate our method’s effectiveness and efficiency in both the qualitative and the quantitative aspects.

2. Related Work

Image manipulation [1–5, 7, 13, 16, 19, 21, 22, 25, 29, 37, 47] aims to modify an input image to a guiding direction. The direction can be a scribble, a mask, a reference image, etc. Recently, text-driven image manipulations [1, 2, 11, 13, 19, 21, 22, 29, 37] have become very popular because textual prompts are intuitive and easily accessible. These methods [2, 11, 20–22, 29] usually utilize a joint image-text semantic space to provide supervision.

One of the famous image-text semantic spaces is CLIP [32], a multi-modal space that contains extremely rich seman-

tics as it is trained using its millions of text-image pairs. It has demonstrated significance on different tasks, such as latent space manipulation [29], domain adaptation [11], and style transfer [22]. Another multi-modal model that has gained much attention is the text-image generative model. Several large-scale text-image models have advanced text-driven image generation, such as Imagen [38], DALL-E2 [34], and latent diffusion [36]. The diffusion model contains the strong mode-capturing ability and the training stability [14, 40]. Some scholars have begun to study how to utilize its powerful capabilities in image manipulation tasks. Most image manipulation methods with diffusion models aim to “hijack” the reverse diffusion process and introduce kinds of guidance and operation [1, 2, 5, 10, 15, 24, 25, 47]. For example, some methods [2, 10, 15, 47] update the intermediate result with the gradient from some conditional models. Some methods also use auxiliary information, such as a mask, to limit the region of the generation [2, 24]. Some method [5] directly replaces the low-frequency information of the intermediate result with that of the reference image and obtains an image with the same structural details. Some method [25] conditions the output image with a stroke image by adding intermediate noises to the image and then denoising it. Some method [20] applies domain adaptation to diffusion models using the CLIP model. In summary, the diffusion process gradually perturbs the data distribution to gaussian noise distribution, while the reverse diffusion process progressively recovers data distribution from the noise.

Meanwhile, another methods [12, 30] uses the diffusion model as prior knowledge in many applications. For example, some scholars [12] mix it with conditional models, e.g., differentiable classifier, and generate specific classes of samples. On the other hand, DreamFusion [30] utilizes differentiable image parameterization [27] and defines this conditional model as a 3D rendering process [26] to generate 3D assets from the text. They can not involve its use in the inference stage [30]. We also take the perspective to use the diffusion model as guidance to avoid the iterative process in the inference stage. Different from utilizing fixed conditional models [12, 30], we define this conditional model as an image-to-image translation model and explore a new scenario where we are optimizing this conditional model. This model, after optimization, can be independent of the diffusion model, thus enabling more efficient image manipulation.

3. Preliminary

3.1. Diffusion Models

Diffusion models are latent-variable generative models that define a Markov chain of diffusion steps to add random noise to data slowly and then learn to reverse the diffusion process to construct desired data samples from the noise

[14, 40].

Suppose the data distribution is $q(x_0)$ where the index 0 denotes the original data. Given a training data sample $x_0 \sim q(x_0)$, the forward diffusion process aims to produce a series of noisy latents x_1, x_2, \dots, x_T by the following Markovian process,

$$q(x_t | x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t \mathbf{I}), \forall t \in [T], \quad (1)$$

where T is the step number of the diffusion process, $[T] = \{1, 2, \dots, T\}$ denotes the set of the index, $\beta_t \in (0, 1)$ represents the variance in the diffusion process, \mathbf{I} is the identity matrix with the same dimensions as the input data x_0 , and $\mathcal{N}(x; \mu, \sigma)$ means the normal distribution with mean μ and covariance σ .

To generate a new data sample, diffusion models sample x_T from the standard normal distribution and then gradually remove noise by the intractable reverse distribution $q(x_{t-1} | x_t)$. Diffusion models learn a neural network p_θ parameterized by θ to approximate the reverse distribution as follows,

$$p_\theta(x_{t-1} | x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)), \quad (2)$$

where μ_θ and Σ_θ are the trainable mean and covariance functions, respectively.

In [14], Σ_θ is simply set as a fixed constant, and μ_θ is reformulated as a function of noise as follows,

$$\mu_\theta(x_t, t) = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{1 - \alpha_t}} \epsilon_\theta(x_t, t) \right), \quad (3)$$

where ϵ_θ is used to predict noise ϵ_t from x_t .

Finally, the diffusion model is trained with simplified evidence lower bound (ELBO) that ignores the weighting term for each timestep t as follows,

$$\mathcal{L}_t(\theta) = \mathbb{E}_{x_0, t, \epsilon} \left[\|\epsilon_t - \epsilon_\theta(x_t, t)\|_2^2 \right]. \quad (4)$$

3.2. Diffusion Models as Prior

According to [12], diffusion models can also be used as off-the-shelf modules in some scenarios, where it $p(x)$ may serve as a prior for another conditional model $c(x, y)$, *i.e.* we can deduce $p(x | y)$ given $p(x)$ and $c(x, y)$. When $c(\mathbf{x}, \mathbf{y})$ is a hard and non-differentiable conditional model, say a deterministic function $x = f(y)$. We optimize y to minimize

$$\sum_t \mathbb{E}_{\epsilon \sim \mathcal{N}(0, \mathbf{I})} \left[\|\epsilon - \epsilon_\theta(\sqrt{\alpha_t}f(y) + \sqrt{1 - \alpha_t}\epsilon, t)\|_2^2 \right]. \quad (5)$$

For example, $f()$ is a latent-variable model that takes latent y as input and generates a sample x . We can regard this equation as sampling y instead of directly sampling images using

¹Please refer supplementary for detailed derivation process.

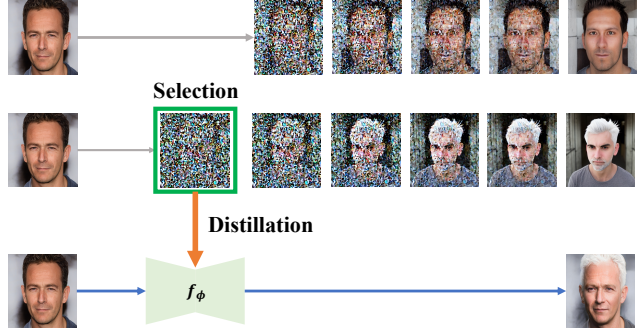


Figure 2. **Core concept of SDD.** Our method involves two steps: 1) selecting the semantically-related timestep and 2) distilling the appropriate semantic knowledge into an image manipulator, f_ϕ .

diffusion models, and inputting y to this conditional model, we will get a sample from the diffusion model. An example of a successful implementation of this idea is DreamFusion [30], where y represents the parameters of a 3D volume and $f()$ represents a volumetric renderer. This method can be used to generate 3D assets from text. In practice, optimizing simultaneously for all t makes it difficult to guide the sample toward a mode. Thus existing methods either anneal t from high to low values [12], or random select t [30]. So the actual optimization process is slightly changed to

$$\mathbb{E}_{\epsilon, t} \left[\|\epsilon - \epsilon_\theta(\sqrt{\alpha_t}f(y) + \sqrt{1 - \alpha_t}\epsilon, t)\|_2^2 \right]. \quad (6)$$

4. Method

In this section, we introduce Selective Diffusion Distillation. The core concept is shown in Fig. 2. First, we introduce how to distill knowledge from a diffusion model into an image manipulation model in Sec. 4.1. Then, we introduce how to select the appropriate timestep in Sec. 4.2.

4.1. Distillation: Learning image manipulator with Diffusion Models

We aim to take the diffusion model as a source of semantic guidance to train a lightweight image manipulator. We first define our problem according to the definition in Sec. 3.2. Training the image manipulator with diffusion prior can also be categorized as an optimization problem with a hard constraint. The hard constraint $f(y)$ becomes the image manipulator in this situation. The difference between our formulation and the previous formulation is the optimized target. We use equation (7) to optimize parameters ϕ of our image manipulator f_ϕ as shown in Fig. 3.

$$\min_{\phi} \mathbb{E}_{\epsilon, t} \left[\|\epsilon - \epsilon_\theta(\sqrt{\alpha_t}f_\phi(y) + \sqrt{1 - \alpha_t}\epsilon, t)\|_2^2 \right]. \quad (7)$$

When optimizing this equation, we follow the approach of skipping the U-Net Jacobian in [30]. The gradient after

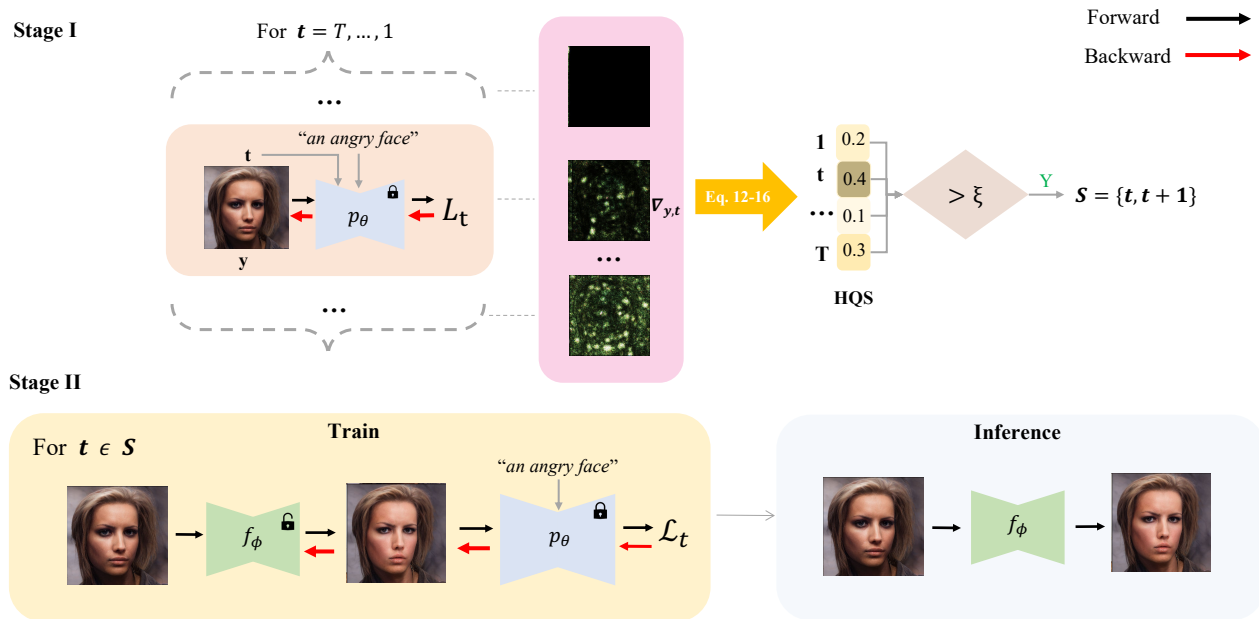


Figure 3. **General framework of Selective Diffusion Distillation (SDD).** **Top: selection stage.** We build an HQS indicator to select appropriate timesteps. **Bottom: distillation stage.** We use the selected timesteps and the pretrain diffusion model p_θ to train the image manipulator f_ϕ .

skipping is equivalent to the noise predicted at the current timestep (based on the diffusion model) minus a random noise. According to the relationship between diffusion models and score-based models [42, 43], this predicted noise from the diffusion model contains the direction from the current distribution to the target distribution. Therefore, we can consider optimizing this equation as guiding our image manipulator by the diffusion model to output results in the target distribution.

This paradigm benefits us from two aspects. Firstly, no matter what semantics we manipulate, we can always ensure fidelity. For general diffusion-based image manipulation, the success of manipulation and the fidelity of the manipulated image are both determined by the noise level, but they conflict with each other. However, for our image manipulator, this conflict does not exist. It is natural to add or tweak regularization terms in the training of our image manipulator so that we can ensure fidelity under different manipulations. Secondly, our method improves inference efficiency. After optimizing the image manipulator, we can perform fast image manipulation through only one forward pass without requiring the diffusion model. Not to mention the manipulator network is lighter than the U-net in the diffusion model. Moreover, our manipulator also demonstrates scalability. Our manipulator finds common knowledge when translating images in the same domain to a specific semantic direction. Once the manipulator is trained, it is easy to reuse this network for manipulating other images without

any retraining. Compared to the efficiency improvement, our extra training cost is little. As described in Sec. 5, when training the manipulator, we only optimize a 4-layer MLP mapper model. This significantly reduces our training costs. When manipulating more images, our method shows faster speed even if we include our extra training time. Detailed quantitative discussion can be found in Sec. 5.3.

4.2. Selection: selecting timestep with the Hybrid Quality Score (HQS)

Assuming fidelity is ensured, the key is to obtain correct semantic guidance from the diffusion model. To achieve this, we need to identify the timestep at which the diffusion model most effectively guides the image manipulator to produce a successful result. We first analyze the gradients provided by the diffusion model. We input the image y into the diffusion model conditioned on textual description γ at every timestep $t \in \{T, \dots, 1\}$, and then compute the gradient on the input image using diffusion training loss:

$$d_t(y, \gamma) = \nabla_y \|\epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t}y + \sqrt{1 - \bar{\alpha}_t}\epsilon, t, \gamma)\|_2^2, \quad (8)$$

where $d_t(y, \gamma) \in \mathbb{R}^{1 \times H \cdot W \cdot C}$ represents the direction that the input image y should move to, given the target distribution γ at the timestep t . If a timestep t is more important, its direction d should have higher quality than other timesteps.

To measure its quality, we have empirically observed that treating this gradient as a confidence map and computing its entropy gives us a good metric. We first convert d_t into

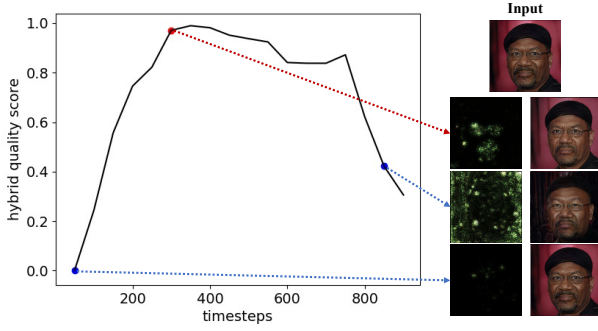


Figure 4. **Effect of our Hybrid Quality Score (HQS).** **Left:** The curve of HQS at different timesteps when the prompt is “angry”. **Right:** Gradient visualization of the corresponding timestep and editing results of training our image manipulator using only this timestep.

a confidence score map p_t through the softmax operation. Each score of the confidence map represents the degree of necessity of the corresponding gradient when modifying the image. Next, for the whole confidence map, we calculated its entropy:

$$H_t = - \sum_i p_t^i \log p_t^i, \quad (9)$$

where p_t^i is the i -th element of p_t . The intuition is that the lower the entropy, the more the $d_t(y, \gamma)$ contains the necessary gradient, so the t is more important. Then, we found this metric susceptible to extreme cases, such as when only a very small part of the gradient map has value. This will result in a very high value of H_t , but the gradient contributes little to the image. Therefore, we introduce the L_1 norm of the gradient $d_t(y, \gamma)$ to avoid this situation:

$$N_t = \sum_i |d_t^i|, \quad (10)$$

where d_t^i is the i -th element of d_t . This metric ensures the magnitude of the overall information of the gradient, thereby avoiding the misjudgment of the entropy metric caused by the local large value. To combine these two metrics, we firstly transform $H = [H_1, \dots, H_T]$ to $\bar{H} = [\bar{H}_1, \dots, \bar{H}_T]$ by using min-max normalization as follows:

$$\bar{H}_t = \frac{H_t - \min(H)}{\max(H) - \min(H)}, \quad (11)$$

and we also do the same to N to get \bar{N} . Significantly, given manipulation target γ , we consider timestep t with lower \bar{H}_t and higher \bar{N}_t to be optimal, so we define a metric called Hybrid Quality Score (HQS):

$$\text{HQS}(\gamma) = \mathbb{E}_y[\bar{N} - \bar{H}], \quad (12)$$

Where $\text{HQS} \in \mathbb{R}^{1 \times T}$. As shown in Fig. 4, training with the timestep with the higher hybrid quality score gives us rich semantic gradients² and better editing results.

Based on this, we propose a timestep selection strategy: Given text prompt γ , we compute the $\text{HQS}(\gamma)$ at each timestep t . Then we build a set of t with a larger HQS value. Next, when optimizing the image manipulator with Eq. (7), we sample timesteps t from the set. To decide the number of timesteps in this set, we introduce a hyperparameter, ξ , which controls the tolerance for uncorrelated t . While selecting t with the maximum HQS is generally effective, relaxing the HQS requirement can also improve editing. By using a smaller ξ , only the most relevant features will change, making the semantic modification relatively simple but requiring fewer optimization costs. On the other hand, using a larger ξ makes the editing more comprehensive but increases the risk of introducing uncorrelated t .

In conclusion, training our image manipulator f_ϕ with this strategy is shown in Alg. 1.

Algorithm 1 Selective Diffusion Distillation

- 1: **Input:** text prompts γ , image data $q(y)$, threshold ξ , pretrained diffusion model ϵ_θ
 - 2: Compute $\text{HQS}(\gamma)$ by Eq. (12)
 - 3: $S = \{t \mid \text{HQS}_t > \xi\}$
 - 4: Randomly initialize our image manipulator f_ϕ
 - 5: **repeat**
 - 6: $y \sim q(y), t \in S, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 - 7: Take gradient descent step on
 - 8: $\nabla_\phi \|\epsilon - \epsilon_\theta(\sqrt{\alpha_t} f_\phi(y) + \sqrt{1 - \alpha_t} \epsilon, t)\|_2^2$
 - 9: **until** converged
-

Using this strategy leads to a set S of a smaller size. This selected S contains fewer ineffective timesteps than the whole timesteps set. Training the image manipulator with selected S leads to correct semantic manipulation. We prove the effectiveness in the Sec. 5.4. The above analysis is agnostic to the concrete form of image manipulators so that it can be extended to any image manipulation framework.

5. Experiments

In this section, we first introduce the implementation details of our Selective Diffusion Distillation. Then, we present the manipulating ability of SDD by showing its application in different domains. Afterward, we compare our method to other image manipulation methods regarding quality and efficiency. At last, we conduct the ablation study of HQS-based step selection to demonstrate its effectiveness.

²we clip and normalize the gradient value for better illustration.

5.1. Implementation details

The discussion in Sec. 4.1 shows that our framework is independent of the form of the image manipulator. Therefore, our method can theoretically be applied to any image manipulation method.

We select StyleGAN [18] as our image manipulator backbone because of its exceptional capabilities in image manipulation [11, 21, 29, 35, 45]. In our framework, both the diffusion model and the StyleGAN contain large parameters. Directly optimizing the StyleGAN by our framework could produce unacceptable computational costs. Therefore, we follow the configuration of [29] to reduce the computational cost. Under this configuration, we only optimize a tiny MLP, called latent mapper, to achieve numerous editing tasks.

For the text-to-image diffusion model, we employ the popular latent diffusion models [36].

Overall, our image manipulator consists of three components: a StyleGAN encoder, a latent mapper, and a StyleGAN generator. The StyleGAN encoder and generator are pre-trained and remain fixed during the optimization. Only the latent mapper is trained in the distillation.

For regularization, we introduce L2 loss and face identity loss [9] as suggested by [29]. We also use gradient clipping in some scenarios for training stability. The whole training procedure follows the Alg. 1.

5.2. Applications

Our SDD shows its strong ability in image manipulation. It successfully edits images of various domains and various attributes while preserving high image quality. Fig. 5 demonstrates our manipulated images. For human faces, we could conduct the hair color, hairstyle, and facial expression transfer, as well as the attributes addition and celebrities conversion. For cats and cars, we could change their types, colors, and some specific details. More manipulation results are shown in our supplementary materials.

5.3. Comparison and Evaluation

In this section, we compare SDD with other image manipulation methods. The empirical results demonstrate our benefits. To measure the result quantitatively, we adopt two metrics, the Fréchet inception distance (FID) [39] and CLIP similarity [33], separately. FID is a metric used to measure the similarity between the distribution of real images and generated images. We use it to measure the similarity between the manipulated and original images. The CLIP similarity measures whether the semantic change of a manipulated image aligns with the text description. A higher directional CLIP similarity indicates a better manipulation result.

Compared to diffusion-based methods We first compare our SDD with other typical diffusion-based image manipula-

tion methods, including SDEdit [25], DDIB [44], and DiffAE [31]. The qualitative is shown in Fig. 6. SDD achieves semantic manipulation in all cases while preserving most of the other information from the input images compared to the baselines.

The quantitative result is shown in Table 1. SDD demonstrates the highest CLIP similarity and maintains the lowest FID, which is consistent with the quantitative comparison. This excellent result indicates that SDD successfully avoids the trade-off problem in diffusion-based manipulation. Moreover, considering that the StyleGAN of our image manipulator is pre-trained with data of the target domain, we also finetune the best of the baselines on this data for a fair comparison. Line 2 of Table 1 shows that despite the extra data, SDD still outperforms it.

In addition, we also compare SDD with these baselines in terms of computation costs. Given the task of manipulating m images for each of n prompts, we compute the total time cost of our method and diffusion-based method to compare their efficiency. For diffusion-based methods, we consider the input image inversion time and the iterative denoising time as their total inference time. $\tau_{\text{Diff, infer}}$ denotes the time cost of each manipulation in diffusion-based methods.

$$\tau_{\text{Diff}} = m \times n \times \tau_{\text{Diff, infer}} \quad (13)$$

For SDD, we need to train individual image manipulators for each of the prompts, and when inference, only the image manipulator is needed. We add the training time of image manipulators and the inference time of manipulation together and treat it as the total time of SDD. $\tau_{\text{SDD, train}}$ denotes the required training time for SDD’s image manipulator to converge. $\tau_{\text{SDD, infer}}$ denotes the time cost of inference with the image manipulator.

$$\tau_{\text{SDD}} = n \times \tau_{\text{SDD, train}} + n \times m \times \tau_{\text{SDD, infer}} \quad (14)$$

According to the Eq. 13 and Eq. 14, all methods have the same complexity of $O(mn)$. However, in our method, the coefficient of mn is $\tau_{\text{SDD, infer}}$, which is a much smaller value compared to $\tau_{\text{Diff, infer}}$.

We further deduce that when m satisfied Eq. (15), our methods will achieve better efficiency. This efficiency improvement will be exaggerated, especially when m becomes larger.

$$m > \frac{\tau_{\text{SDD, train}}}{\tau_{\text{Diff, infer}} - \tau_{\text{SDD, infer}}} \quad (15)$$

We compared the overall time cost between our method and diffusion-based methods in Table 1. All diffusion models use DDIM [41] acceleration with 50 inference steps. We set $m=100$, and $n=10$ in the experiment. The comparison of computational costs is shown in column 3 of Table 1.



Figure 5. The manipulation results of SDD of various domains (CelebA-HQ [23], AFHQ-cat [8], LSUN-car [46]) and various attributes. We keep the image fidelity and make it coherent with the text.

Compared to StyleCLIP Our method shares some similarities to the StyleCLIP in the aspect of the manipulator. Here, we discuss the major difference between SDD and StyleCLIP and empirically compare them. The diffusion model has a significant advantage over CLIP in that the gradient it provides shares the same spatial size as images. Compared to CLIP, these gradients from diffusion models contain structural information, making our SDD capable of position manipulations. In contrast, CLIP guidance is insensitive to 3D positional information, so it does not support

such manipulation. Other studies [17] also notice this insensitivity of CLIP, and they provide further explanation for that. Concretely, our image manipulator can change the pose of a human’s face, but CLIP-based methods fails, as shown in Fig. 7.

We also quantitatively compare manipulation tasks that can be accomplished by both methods and demonstrate that our SDD provides better guidance for training the image manipulator, as shown in Table 2. Note that directly replacing CLIP with a diffusion model and not using our timestep selection strategy will not yield such results.

	FID	CLIP Similarity	Total time
SDEdit	32.126	0.2189	2215
SDEdit*	16.761	0.2133	2215
DDIB	87.737	0.2268	3502
DiffAE	41.896	0.2136	5658
SDD	6.066	0.2337	148.67

Table 1. The quantitative comparison between our method and diffusion-based image manipulations. * means we fine-tune the diffusion model on the target dataset.

5.4. Ablation Study

In this section, we conduct an ablation study of the timestep selection strategy for optimizing the image manipulator. We build four configurations: For random strategy, we randomly sample t from all timesteps; For small threshold strategy, we use HQS with a small ξ to obtain S with a large number of t ; For large threshold strategy, we use HQS with a large ξ to obtain S with a small number of t ; For



Figure 6. **The visual comparison between SDD and diffusion-based image manipulations.** α ranges from zero to one and represents the noise level. SDEdit and DDIB fail to manipulate the semantics at a noise level of 0.7, while SDD successfully manipulates the semantics with less distortion.

	Domains	FID	CLIP Similarity
StyleCLIP	Face	16.542	0.2250
	Car	52.356	0.2652
	Cat	42.737	0.2582
SDD	Face	6.066	0.2337
	Car	52.356	0.2621
	Cat	39.354	0.2948

Table 2. Quantitative comparison between SDD and StyleCLIP.



Figure 7. **Comparison of SDD and StyleCLIP in pose manipulation.** Text prompt: “side view”. The results demonstrate that the diffusion model provides superior semantic guidance, enabling a broader range of manipulations.

largest HQS strategy, we only sample t with the largest HQS value. We keep other configurations the same and compare them qualitatively and quantitatively.

The qualitative result in Fig. 8 shows that the proposed HQS-based step selection significantly overperforms other baseline methods. Random strategy always seems to have little modification to the original image. We attribute this




Figure 8. **Visual result of Ablation Study.** The average HQS score of timesteps used for training increases from left to right and so does the accuracy of manipulation.

to its redundant timestep selection. The result of the largest HQS strategy is the most desirable and intensive modification, proving that our HQS helps us find the most useful step. Furthermore, by combining the results from the small threshold, large threshold, and largest HQS, we can observe that under the same number of training iterations, the average HQS score of the t they sampled increased in order, leading to a sequential improvement of the manipulation effect. Therefore, it is proved that HQS can select the t with the maximum contribution to the semantic manipulation. The quantitative result in Table 3 also demonstrates that the largest HQS strategy performs the best. The increase in FID is caused by the manipulation effect, as shown in Fig. 8. Meanwhile, we still preserve image quality, as evidenced by our very low FID.

	FID	CLIP Similarity
Random	3.288	0.2146
Small threshold	3.819	0.2155
Large threshold	5.927	0.2168
Largest HQS	9.154	0.2190

Table 3. The quantitative ablation for HQS-based timestep selection strategy. The FID increase because the manipulation effect increase, as shown in Fig. 8

6. Conclusion

In this paper, we present a novel image manipulation method called Selective Distillation Diffusion (SDD). This paradigm avoids the Editability & Fidelity trade-off by distilling the diffusion model’s knowledge to a lightweight image manipulator. To distillate correct semantic information, we carefully design the Hybrid Quality Score (HQS) to help us select the helpful timesteps. We evaluate our method SDD on a variety of image manipulation tasks and achieve promising results.

7. Acknowledgement

This work is supported by National Natural Science Foundation of China (No. 62206068).

References

- [1] Omri Avrahami, Ohad Fried, and Dani Lischinski. Blended latent diffusion. *arXiv preprint arXiv:2206.02779*, 2022. 1, 2
- [2] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. In *CVPR*, 2022. 1, 2
- [3] Ying-Cong Chen, Xiaogang Xu, and Jiaya Jia. Domain adaptive image-to-image translation. In *CVPR*, pages 5274–5283, 2020. 2
- [4] Ying-Cong Chen, Xiaogang Xu, Zhuotao Tian, and Jiaya Jia. Homomorphic latent space interpolation for unpaired image-to-image translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2408–2416, 2019. 2
- [5] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. In *ICCV*, 2021. 2
- [6] Jooyoung Choi, Jungbeom Lee, Chaehun Shin, Sungwon Kim, Hyunwoo Kim, and Sungroh Yoon. Perception prioritized training of diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11472–11481, 2022. 1
- [7] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8789–8797, 2018. 2
- [8] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *CVPR*, 2020. 7
- [9] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019. 6
- [10] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *NeurIPS*, 2021. 1, 2
- [11] Rinon Gal, Or Patashnik, Haggai Maron, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. Stylegan-nada: Clip-guided domain adaptation of image generators. *ACM Transactions on Graphics (TOG)*, 2022. 2, 6
- [12] Alexandros Graikos, Nikolay Malkin, Nebojsa Jojic, and Dimitris Samaras. Diffusion models as plug-and-play priors. *arXiv preprint arXiv:2206.09012*, 2022. 2, 3
- [13] 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
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 2020. 1, 2, 3
- [15] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 2
- [16] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. 2
- [17] Ajay Jain, Matthew Tancik, and Pieter Abbeel. Putting nerf on a diet: Semantically consistent few-shot view synthesis. In *ICCV*, 2021. 7
- [18] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019. 6
- [19] 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. *arXiv preprint arXiv:2210.09276*, 2022. 2
- [20] Gwanghyun Kim, Taesung Kwon, and Jong Chul Ye. Diffusionclip: Text-guided diffusion models for robust image manipulation. In *CVPR*, 2022. 2
- [21] Umut Kocasari, Alara Dirik, Mert Tiftikci, and Pinar Yarnadag. Stylemc: Multi-channel based fast text-guided image generation and manipulation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 895–904, 2022. 2, 6
- [22] Gihyun Kwon and Jong Chul Ye. Clipstyler: Image style transfer with a single text condition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18062–18071, 2022. 2
- [23] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, December 2015. 7
- [24] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *CVPR*, pages 11461–11471, 2022. 1, 2
- [25] Chenlin Meng, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021. 1, 2, 6
- [26] B Mildenhall, PP Srinivasan, M Tancik, JT Barron, R Ramamoorthi, and R Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 2
- [27] Alexander Mordvintsev, Nicola Pezzotti, Ludwig Schubert, and Chris Olah. Differentiable image parameterizations. *Distill*, 2018. 2
- [28] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. 1
- [29] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *CVPR*, pages 2085–2094, 2021. 2, 6
- [30] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 2, 3
- [31] Konpat Preechakul, Nattanat Chatthee, Suttisak Wizadwongsa, and Supasorn Suwajanakorn. Diffusion autoencoders: Toward a meaningful and decodable representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10619–10629, 2022. 6
- [32] 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*, 2021. 2

- [33] 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 *ICML*, 2021. 6
- [34] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 1, 2
- [35] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *CVPR*, 2021. 6
- [36] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 1, 2, 6
- [37] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. *arXiv preprint arXiv:2208.12242*, 2022. 2
- [38] 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. *arXiv preprint arXiv:2205.11487*, 2022. 1, 2
- [39] Maximilian Seitzer. pytorch-fid: FID Score for PyTorch. <https://github.com/mseitzer/pytorch-fid>, August 2020. Version 0.3.0. 6
- [40] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. 2015. 1, 2, 3
- [41] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 1, 6
- [42] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019. 1, 4
- [43] 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. 1, 4
- [44] Xuan Su, Jiaming Song, Chenlin Meng, and Stefano Ermon. Dual diffusion implicit bridges for image-to-image translation. *arXiv preprint arXiv:2203.08382*, 2022. 1, 6
- [45] Zongze Wu, Dani Lischinski, and Eli Shechtman. Stylespace analysis: Disentangled controls for stylegan image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 6
- [46] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv:1506.03365*, 2015. 7
- [47] Min Zhao, Fan Bao, Chongxuan Li, and Jun Zhu. Egsde: Unpaired image-to-image translation via energy-guided stochastic differential equations. *arXiv preprint arXiv:2207.06635*, 2022. 2